

HP 10Gb Ethernet BL-c Switch

Application Guide



Part number: 445946-001

First edition: June 2007

Legal notices

© 2007 Hewlett-Packard Development Company, L.P.

The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Microsoft®, Windows®, and Windows NT® are U.S. registered trademarks of Microsoft Corporation.

SunOS™ and Solaris™ are trademarks of Sun Microsystems, Inc. in the U.S. and other countries.

Cisco® is a registered trademark of Cisco Systems, Inc. and/or its affiliates in the U.S. and certain other countries.

Contents

Accessing the switch

Introduction	9
Additional references	10
Typographical conventions	10
Management Network	10
Connecting through the console port	11
Connecting through Telnet	11
Connecting through Secure Shell	11
Using the command line interfaces	12
Configuring an IP interface	12
Using the Browser-based Interface	13
Using Simple Network Management Protocol	14
SNMP v1.0	14
SNMP v3.0	14
Default configuration	14
User configuration	15
View based configurations	16
CLI user equivalent	16
CLI oper equivalent	17
Configuring SNMP trap hosts	17
SNMPv1 trap host	17
SNMPv2 trap host configuration	19
SNMPv3 trap host configuration	19
Secure access to the switch	20
Setting allowable source IP address ranges	20
Configuring an IP address range for the management network	21
RADIUS authentication and authorization	21
How RADIUS authentication works	21
Configuring RADIUS on the switch (CLI example)	22
Configuring RADIUS on the switch (BBI example)	23
RADIUS authentication features	24
User accounts for RADIUS users	24
RADIUS attributes for user privileges	25
TACACS+ authentication	25
How TACACS+ authentication works	26
TACACS+ authentication features	26
Authorization	26
Accounting	27
Configuring TACACS+ authentication on the switch (CLI example)	28
Configuring TACACS+ authentication on the switch (BBI example)	29
Secure Shell and Secure Copy	30
Configuring SSH and SCP features (CLI example)	31
Using SSH and SCP client commands	32
SSH and SCP encryption of management messages	33
Generating RSA host and server keys for SSH access	33
SSH/SCP integration with RADIUS and TACACS+ authentication	34

User access control	34
Setting up user IDs	35
Ports and trunking	
Introduction	36
Ports on the switch	36
Port trunk groups.....	37
Statistical load distribution.....	37
Built-in fault tolerance.....	37
Before you configure trunks	37
Trunk group configuration rules.....	38
Port trunking example	39
Configuring trunk groups (CLI example)	40
Configuring trunk groups (BBI example)	41
Configurable Trunk Hash algorithm	44
Link Aggregation Control Protocol	44
Configuring LACP	46
Port-based Network Access and traffic control	
Port-based Network Access control.....	47
Extensible authentication protocol over LAN	47
802.1x authentication process.....	48
EAPoL Message Exchange	48
802.1x port states.....	49
Supported RADIUS attributes	50
EAPoL configuration guidelines	51
Port-based traffic control	51
Configuring port-based traffic control.....	52
VLANs	
Introduction	53
Overview	53
VLANs and port VLAN ID numbers	53
VLAN numbers	53
Viewing VLANs	54
PVID numbers	54
Viewing and configuring PVIDs.....	54
Port information	54
Port configuration.....	54
VLAN tagging	55
VLANs and IP interfaces.....	58
VLAN topologies and design considerations.....	58
VLAN configuration rules	59
Multiple VLANs with tagging.....	60
Configuring the example network.....	61
Configuring ports and VLANs on Switch 1 (CLI example)	61
Configuring ports and VLANs on Switch 2 (CLI example)	63
Configuring ports and VLANs on Switch 1 (BBI example)	64
FDB static entries.....	66
Trunking support for FDB static entries.....	67
Configuring a static FDB entry	67
Spanning Tree Protocol	

Introduction	68
Overview	68
Bridge Protocol Data Units	68
Determining the path for forwarding BPDUs	69
Bridge priority	69
Port priority	69
Port path cost	69
Spanning Tree Group configuration guidelines	69
Default Spanning Tree configuration	69
Adding a VLAN to a Spanning Tree Group	70
Creating a VLAN	70
Rules for VLAN tagged ports	70
Adding and removing ports from STGs	70
Assigning cost to ports and trunk groups	71
Multiple Spanning Trees	71
Why do we need Multiple Spanning Trees?	71
VLAN participation in Spanning Tree Groups	72
Configuring Multiple Spanning Tree Groups	73
Configuring Switch 1 (CLI example)	73
Configuring Switch 2 (CLI example)	73
Configuring Switch 1 (BBI example)	74
Port Fast Forwarding	76
Configuring Port Fast Forwarding	76
Fast Uplink Convergence	76
Configuration guidelines	76
Configuring Fast Uplink Convergence	76
RSTP and MSTP	
Introduction	77
Rapid Spanning Tree Protocol	77
Port state changes	77
Port type and link type	78
Edge port	78
Link type	78
RSTP configuration guidelines	78
RSTP configuration example	78
Configuring Rapid Spanning Tree (CLI example)	78
Configuring Rapid Spanning Tree Protocol (BBI example)	79
Multiple Spanning Tree Protocol	80
MSTP region	80
Common Internal Spanning Tree	80
MSTP configuration guidelines	81
MSTP configuration example	81
Configuring Multiple Spanning Tree Protocol (CLI example)	81
Configuring Multiple Spanning Tree Protocol (BBI example)	82
Quality of Service	
Introduction	86
Overview	86
Using ACL filters	87
Summary of packet classifiers	87
Summary of ACL actions	89
Understanding ACL precedence	89

Using ACL Groups	90
ACL Metering and Re-marking	91
Metering	91
Re-marking	91
Viewing ACL statistics	91
ACL configuration examples	92
Configure Access Control Lists (CLI example)	92
Configure Access Control Lists and Groups (BBI example 1)	93
Using DSCP values to provide QoS	97
Differentiated Services concepts	97
Per Hop Behavior	97
QoS levels	98
Using 802.1p priorities to provide QoS	98
802.1p configuration (CLI example)	100
802.1p configuration (BBI example)	100
Queuing and scheduling	105
Basic IP routing	
IP routing benefits	106
Routing between IP subnets	106
Example of subnet routing	109
Using VLANs to segregate broadcast domains	110
Dynamic Host Configuration Protocol	112
DHCP relay agent	112
DHCP relay agent configuration	113
Routing Information Protocol	
Distance vector protocol	114
Stability	114
Routing updates	114
RIPv1	115
RIPv2	115
RIPv2 in RIPv1 compatibility mode	115
RIP Features	115
Poison	115
Triggered updates	115
Multicast	116
Default	116
Metric	116
Authentication	116
RIP configuration example	117
IGMP Snooping	
Introduction	118
Overview	118
IGMPv3	119
FastLeave	119
IGMP Filtering	120
Configuring the range	120
Configuring the action	120
Static multicast router	121
IGMP Snooping configuration example	121

Configuring IGMP Snooping (CLI example)	121
Configuring IGMP Filtering (CLI example).....	122
Configuring a Static Mrouter (CLI example).....	122
Configuring IGMP Snooping (BBI example).....	123
Configuring IGMP Filtering (BBI example)	125
Configuring a Static Multicast Router (BBI example).....	129

OSPF

OSPF overview	131
Types of OSPF areas	131
Types of OSPF routing devices	132
Neighbors and adjacencies	133
Link-State Database	133
Shortest Path First Tree	133
Internal versus external routing.....	134
OSPF implementation in HP 10GbE switch software	134
Configurable parameters	134
Defining areas.....	135
Assigning the area index	135
Using the area ID to assign the OSPF area number	136
Attaching an area to a network	136
Interface cost.....	136
Electing the designated router and backup	137
Summarizing routes.....	137
Default routes	137
Virtual links	138
Router ID.....	138
Authentication	139
Host routes for load balancing.....	140
OSPF features not supported in this release	141
OSPF configuration examples.....	141
Example 1: Simple OSPF domain (CLI example).....	141
Example 1: Simple OSPF domain (BBI example)	142
Example 2: Virtual links	150
Configuring OSPF for a virtual link on Switch A.....	150
Configuring OSPF for a virtual link on Switch B.....	151
Other Virtual Link Options.....	152
Example 3: Summarizing routes.....	152
Verifying OSPF configuration.....	154

Remote monitoring

Introduction	155
Overview	155
RMON group 1—statistics	155
Configuring RMON Statistics (CLI example).....	156
Configuring RMON Statistics (BBI example).....	156
RMON group 2—history.....	158
History MIB objects	159
RMON group 3—alarms	161
Alarm MIB objects.....	161
RMON group 9—events	165
Configuring RMON Events (CLI example).....	165
Configuring RMON Events (BBI example).....	166

High availability

Introduction	167
Uplink Failure Detection	167
Failure Detection Pair.....	168
Spanning Tree Protocol with UFD	168
Configuration guidelines	169
Monitoring Uplink Failure Detection.....	169
Configuring Uplink Failure Detection.....	169
Configuring UFD on Switch 1 (CLI example)	170
Configuring UFD on Switch 2 (CLI example)	170
Configuring Uplink Failure Detection (BBI example)	171
VRRP overview.....	173
VRRP components.....	173
Virtual router	173
Virtual router MAC address	173
Owners and renters.....	173
Master and backup virtual router	174
Virtual Interface Router.....	174
VRRP operation.....	174
Selecting the master VRRP router	174
Failover methods.....	175
Active-Active redundancy	175
HP 10GbE switch extensions to VRRP	176
Tracking VRRP router priority	176
Virtual router deployment considerations	177
Assigning VRRP virtual router ID	177
Configuring the switch for tracking	177
High availability configurations.....	178
Active-Active configuration	178
Task 1: Configure Switch A.....	178
Task 2: Configure Switch B	180
Task 1: Configure Switch A (BBI example).....	181

Troubleshooting tools

Introduction	191
Port Mirroring.....	191
Configuring Port Mirroring (CLI example)	192
Configuring Port Mirroring (BBI example)	193
Other network troubleshooting techniques	195
Console and Syslog messages	195
Ping	195
Trace route.....	195
Statistics and state information	195
Customer support tools.....	195

Index

Accessing the switch

Introduction

This guide will help you plan, implement, and administer the switch software for the HP 10Gb Ethernet BL-c Switch. Where possible, each section provides feature overviews, usage examples, and configuration instructions.

- “Accessing the switch” describes how to configure and view information and statistics on the switch over an IP network. This chapter also discusses different methods to manage the switch for remote administrators, such as setting specific IP addresses and using Remote Authentication Dial-in User Service (RADIUS) authentication, Secure Shell (SSH), and Secure Copy (SCP) for secure access to the switch.
- “Ports and port trunking” describes how to group multiple physical ports together to aggregate the bandwidth between large-scale network devices.
- “Port-based Network Access and Traffic Control” describes how to authenticate devices attached to a LAN port that has point-to-point connection characteristics. Port-based Network Access Control provides security to ports of the HP 10GbE switch that connect to servers. Port-based Traffic Control allows the switch to guard against broadcast storms.
- “VLANs” describes how to configure Virtual Local Area Networks (VLANs) for creating separate network segments, including how to use VLAN tagging for devices that use multiple VLANs.
- “Spanning Tree Protocol” discusses how spanning trees configure the network so that the switch uses the most efficient path when multiple paths exist.
- “Rapid Spanning Tree Protocol/Multiple Spanning Tree Protocol” describes extensions to the Spanning Tree Protocol that provide rapid convergence of spanning trees for fast reconfiguration of the network.
- “Quality of Service” discusses Quality of Service features, including IP filtering using Access Control Lists, Differentiated Services, and IEEE 802.1p priority values.
- “Basic IP Routing” describes how to configure the HP 10GbE switch for IP routing using IP subnets, and DHCP Relay.
- “Routing Information Protocol” describes how the HP 10GbE switch software implements standard Routing Information Protocol (RIP) for exchanging TCP/IP route information with other routers.
- “IGMP Snooping” describes how to use IGMP to conserve bandwidth in a multicast-switching environment.
- “OSPF” describes OSPF concepts, how OSPF is implemented, and examples of how to configure your switch for OSPF support.
- “Remote Monitoring” describes how to configure the RMON agent on the switch, so the switch can exchange network monitoring data.
- “High Availability” describes how the HP 10GbE switch supports high-availability network topologies. This release provides Uplink Failure Detection and Virtual Router Redundancy Protocol (VRRP).
- “Troubleshooting tools” describes Port Mirroring and other troubleshooting techniques.

Additional references

Additional information about installing and configuring the switch is available in the following guides, which are available at <http://www.hp.com/go/bladefsystem/documentation>.

- *HP 10Gb Ethernet BL-c Switch User Guide*
- *HP 10Gb Ethernet BL-c Switch Command Reference Guide*
- *HP 10Gb Ethernet BL-c Switch ISCLI Reference Guide*
- *HP 10Gb Ethernet BL-c Switch Browser-based Interface Reference Guide*
- *HP 10Gb Ethernet BL-c Switch Quick Setup Instructions*

Typographical conventions

The following table describes the typographic styles used in this guide:

Table 1 Typographic conventions

Typeface or symbol	Meaning	Example
AaBbCc123	This type depicts onscreen computer output and prompts.	Main#
AaBbCc123	This type displays in command examples and shows text that must be typed in exactly as shown.	Main# sys
<AaBbCc123>	This bracketed type displays in command examples as a parameter placeholder. Replace the indicated text with the appropriate real name or value when using the command. Do not type the brackets.	To establish a Telnet session, enter: host# telnet <IP address> Read your user guide thoroughly.
[]	Command items shown inside brackets are optional and can be used or excluded as the situation demands. Do not type the brackets.	host# ls [-a]

Management Network

The 10Gb Ethernet BL-c Switch is an integral subsystem within the overall BladeSystem. The BladeSystem chassis includes an Onboard Administrator as the central element for overall chassis and control.

The 10GbE switch communicates with the Onboard Administrator through its internal management port (port 17). The factory default settings permit management and control access to the switch through the 10/100 Mbps Ethernet port on the Onboard Administrator, or the built-in console port. You also can use the external Ethernet ports to manage and control the 10GbE switch.

The 10GbE switch management network has the following characteristics:

- Port 17—Management port 17 has a fixed configuration, as follows:
 - 100 Mbps
 - Full duplex
 - Flow control: both
 - No auto-negotiation

- Untagged
- Port VLAN ID (PVID): 4095
- VLAN 4095—Management VLAN 4095 isolates management traffic within the HP 10GbE switch. VLAN 4095 contains only one member port (port 17). No other ports can be members of VLAN 4095.
- Interface 250—Management interface 250 is associated with VLAN 4095. No other interfaces can be associated with VLAN 4095. You can configure the IP address of the management interface manually or through Dynamic Host Control Protocol (DHCP).
- Gateway 254—This gateway is the default gateway for the management interface.
- STG 128—If the HP 10GbE switch is configured to use multiple spanning trees, spanning tree group 128 (STG 128) contains management VLAN 4095, and no other VLANs are allowed in STG 128. The default status of STG 128 is `off`.

If the 10GbE switch is configured to use Rapid Spanning Tree Protocol, STG 1 contains management VLAN 4095.

To access the 10GbE switch management interface through the Onboard Administrator:

- Use the Onboard Administrator internal DHCP server, through Enclosure-Based IP Addressing
- Use an external DHCP server. Connect the Onboard administrator and the 10GbE switch to the network, and disable Enclosure-Based IP Addressing.
- Assign a static IP interface to the Onboard Administrator and to the 10GbE switch management interface (interface 250).

Connecting through the console port

Using a null modem cable, you can directly connect to the switch through the console port. A console connection is required in order to configure Telnet or other remote access applications. For more information on establishing console connectivity to the switch, see the *HP 10Gb Ethernet BL-c Switch User Guide*.

Connecting through Telnet

By default, Telnet is enabled on the switch. Once the IP parameters are configured, you can access the CLI from any workstation connected to the network using a Telnet connection. Telnet access provides the same options for a user and an administrator as those available through the console port, minus certain commands. The switch supports four concurrent Telnet connections.

To establish a Telnet connection with the switch, run the Telnet program on your workstation and issue the `telnet` command, followed by the switch IP address:

```
telnet <switch IP address>
```

Connecting through Secure Shell

By default, the Secure Shell (SSH) protocol is disabled on the switch. SSH enables you to securely log into another computer over a network to execute commands remotely. As a secure alternative to using Telnet to manage switch configuration, SSH ensures that all data sent over the network is encrypted and secure. For more information, see the “Secure Shell and Secure Copy” section later in this chapter. For additional information on the CLI, see the *HP 10Gb Ethernet BL-c Switch Command Reference Guide*.

Using the command line interfaces

The command line interface (CLI) can be accessed via local terminal connection or a remote session using Telnet or SSH. The CLI is the most direct method for collecting switch information and performing switch configuration.

The HP 10GbE switch provides two CLI modes: The menu-based AOS CLI, and the tree-based ISCLI. You can set the HP 10GbE switch to use either CLI mode.

The Main Menu of the AOS CLI, with administrator privileges, is displayed below:

```
[Main Menu]
  info      - Information Menu
  stats     - Statistics Menu
  cfg       - Configuration Menu
  oper      - Operations Command Menu
  boot      - Boot Options Menu
  maint     - Maintenance Menu
  diff      - Show pending config changes [global command]
  apply     - Apply pending config changes [global command]
  save      - Save updated config to FLASH [global command]
  revert    - Revert pending or applied changes [global command]
  exit      - Exit [global command, always available]
```

For complete information about the AOS CLI, see the *HP 10Gb Ethernet BL-c Switch Command Reference Guide*.

The ISCLI provides a tree-based command structure, for users familiar with similar products.

An example of a typical ISCLI command is displayed below:

```
Switch(config)# spanning-tree stp 1 enable
```

For complete information about the ISCLI, refer to the *ISCLI Reference Guide*.

Configuring an IP interface

An IP interface address must be set on the switch to provide management access to the switch over an IP network. By default, the management interface is set up to request its IP address from a Bootstrap Protocol (BOOTP) server.

If you have a BOOTP server on your network, add the Media Access Control (MAC) address of the switch to the BOOTP configuration file located on the BOOTP server. The MAC address can be found on a small white label on the back panel of the switch. The MAC address can also be found in the System Information menu (see the *HP 10Gb Ethernet BL-c Switch Command Reference Guide* or *ISCLI Guide*.) If you are using a DHCP server that also does BOOTP, you do not have to configure the MAC address.

If you do not have a BOOTP server, you must manually configure an IP address.

The following example shows how to manually configure an IP address on the switch:

1. Configure an IP interface for the Telnet connection, using the sample IP address of 205.21.17.3.
2. The pending subnet mask address and broadcast address are automatically calculated.

```
>> # /cfg/l3/if 1 (Select IP interface 1)
>> IP Interface 1# addr 205.21.17.3 (Assign IP address for the interface)
Current IP address: 0.0.0.0
New pending IP address: 205.21.17.3
Pending new subnet mask: 255.255.255.0
. . . . .
>> IP Interface 1# ena (Enable IP interface 1)
```

3. If necessary, configure up to two default gateways.
4. Configuring the default gateways allows the switch to send outbound traffic to the routers.

```
>> IP Interface 5# ../gw 1 (Select primary default gateway)
>> Default gateway 1# addr 205.21.17.1 (Assign IP address for primary router)
>> Default gateway 1# ena (Enable primary default gateway)
>> Default gateway 1# ../gw 2 (Select secondary default gateway)
>> Default gateway 2# addr 205.21.17.2 (Assign address for secondary router)
>> Default gateway 2# ena (Enable secondary default gateway)
```

5. Apply, verify, and save the configuration.

```
>> Default gateway 2# apply (Apply the configuration)
>> Default gateway 2# save (Save the configuration)
>> # /cfg/dump (Verify the configuration)
```

Using the Browser-based Interface

By default, the Browser-based Interface (BBI) protocol is enabled on the switch. The Browser-based Interface (BBI) provides access to the common configuration, management and operation features of the switch through your Web browser. For more information, see the *HP 10Gb Ethernet BL-c Switch Browser-based Interface Reference Guide*.

The BBI is organized at a high level as follows:

- Configuration—These menus provide access to the configuration elements for the entire switch.
 - System—Configure general switch configuration elements.
 - Switch ports—Configure switch ports and related features.
 - Port-based port mirroring—Configure mirrored ports and monitoring ports.
 - Layer 2—Configure Layer 2 features, including trunk groups, VLANs, and Spanning Tree Protocol.
 - RMON menu—Configure Remote Monitoring (RMON) functions.
 - Layer 3—Configure all of the IP related information, including IGMP Snooping.
 - QoS—Configure Quality of Service features.
 - Access Control—Configure Access Control Lists and Groups.
 - Uplink Failure Detection—Configure a Failover Pair of Links to Monitor and Links to Disable.
- Statistics—These menus provide access to the switch statistics and state information.
- Dashboard—These menus display settings and operating status of a variety of switch features.

Using Simple Network Management Protocol

The switch software provides SNMP v1.0 and SNMP v3.0 support for access through any network management software, such as HP-OpenView.

SNMP v1.0

To access the SNMP agent on the switch, the read and write community strings on the SNMP manager should be configured to match those on the switch. The default read community string on the switch is `public` and the default write community string is `private`.

The read and write community strings on the switch can be changed using the following commands on the CLI.

```
>> /cfg/sys/ssnmp/rcomm
```

-and-

```
>> /cfg/sys/ssnmp/wcomm
```

The SNMP manager should be able to reach the management interface or any one of the IP interfaces on the switch.

For the SNMP manager to receive the traps sent out by the SNMP agent on the switch, the trap host on the switch should be configured with the following command:

```
/cfg/sys/ssnmp/snmpv3/taddr
```

For more details, see “Configuring SNMP trap hosts”.

SNMP v3.0

SNMPv3 is an enhanced version of the Simple Network Management Protocol, approved by the Internet Engineering Steering Group in March, 2002. SNMP v3.0 contains additional security and authentication features that provide data origin authentication, data integrity checks, timeliness indicators, and encryption to protect against threats such as masquerade, modification of information, message stream modification, and disclosure.

SNMP v3 ensures that the client can use SNMP v3 to query the MIBs, mainly for security.

To access the SNMP v3.0 menu, enter the following command in the CLI:

```
>> # /cfg/sys/ssnmp/snmpv3
```

For more information on SNMP MIBs and the commands used to configure SNMP on the switch, see the *HP 10Gb Ethernet BL-c Switch Command Reference Guide*.

Default configuration

The switch software has two users by default. Both the users `adminmd5` and `adminsha` have access to all the MIBs supported by the switch.

- username 1—`adminmd5/password adminmd5`. Authentication used is MD5.
- username 2—`adminsha/password adminsha`. Authentication used is SHA.
- username 3—`v1v2only/password none`.

To configure an SNMP user name, enter the following command from the CLI:

```
>> # /cfg/sys/ssnmp/snmpv3/usm 6
```

User configuration

Users can be configured to use the authentication/privacy options. The HP 10GbE switch supports two authentication algorithms: MD5 and SHA, as specified in the following command:

```
/cfg/sys/ssnmp/snmpv3/usm <x>/auth md5|sha
```

1. To configure a user with name `admin`, authentication type MD5, authentication password of `admin`, and privacy option DES with privacy password of `admin`, use the following CLI commands:

```
>> # /cfg/sys/ssnmp/snmpv3/usm 5
>> SNMPv3 usmUser 5 # name "admin" (Configure 'admin' user type)
>> SNMPv3 usmUser 5 # auth md5
>> SNMPv3 usmUser 5 # authpw admin
>> SNMPv3 usmUser 5 # priv des
>> SNMPv3 usmUser 5 # privpw admin
```

2. Configure a user access group, along with the views the group may access. Use the access table to configure the group's access level.

```
>> # /cfg/sys/ssnmp/snmpv3/access 5
>> SNMPv3 vacmAccess 5 # name "admingrp" (Configure an access group)
>> SNMPv3 vacmAccess 5 # level authPriv
>> SNMPv3 vacmAccess 5 # rview "iso"
>> SNMPv3 vacmAccess 5 # wview "iso"
>> SNMPv3 vacmAccess 5 # nview "iso"
```

Because the read view (`rview`), write view (`wview`), and notify view (`nview`) are all set to "iso," the user type has access to all private and public MIBs.

3. Assign the user to the user group. Use the group table to link the user to a particular access group.

```
>> # /cfg/sys/ssnmp/snmpv3/group 5
>> SNMPv3 vacmSecurityToGroup 5 # uname admin
>> SNMPv3 vacmSecurityToGroup 5 # gname admingrp
```

If you want to allow user access only to certain MIBs, see the "View based configurations" section.

View based configurations

CLI user equivalent

To configure an SNMP user equivalent to the CLI user, use the following configuration:

```
/c/sys/ssnmp/snmpv3/usm 4
name "usr" (Configure the user)
/c/sys/ssnmp/snmpv3/access 3
name "usrgrp" (Configure access group 3)
rview "usr"
wview "usr"
nview "usr"
/c/sys/ssnmp/snmpv3/group 4 (Assign user to access group 3)
uname usr
gname usrgrp
/c/sys/ssnmp/snmpv3/view 6 (Create views for user)
name "usr"
tree " 1.3.6.1.4.1.11.2.3.7.11.33.1.2.1.2" (Agent statistics)
/c/sys/ssnmp/snmpv3/view 7
name "usr"
tree " 1.3.6.1.4.1.11.2.3.7.11.33.1.2.1.3" (Agent information)
/c/sys/ssnmp/snmpv3/view 8
name "usr"
tree " 1.3.6.1.4.1.11.2.3.7.11.33.1.2.2.2" (L2 statistics)
/c/sys/ssnmp/snmpv3/view 9
name "usr"
tree " 1.3.6.1.4.1.11.2.3.7.11.33.1.2.2.3" (L2 information)
/c/sys/ssnmp/snmpv3/view 10
name "usr"
tree " 1.3.6.1.4.1.11.2.3.7.11.33.1.2.3.2" (L3 statistics)
/c/sys/ssnmp/snmpv3/view 11
name "usr"
tree " 1.3.6.1.4.1.11.2.3.7.11.33.1.2.3.3" (L3 information)
```


CLI oper equivalent

To configure an SNMP user equivalent to the CLI `oper`, use the following configuration:

```

/c/sys/ssnmp/snmpv3/usm 5
name "oper" (Configure the oper)
/c/sys/ssnmp/snmpv3/access 4
name "opergrp" (Configure access group 4)
rview "oper"
wview "oper"
nview "oper"
/c/sys/ssnmp/snmpv3/group 4 (Assign oper to access group 4)
uname oper
gname opergrp
/c/sys/ssnmp/snmpv3/view 20 (Create views for oper)
name "oper"
tree " 1.3.6.1.4.1.11.2.3.7.11.33.1.2.1.2" (Agent statistics)
/c/sys/ssnmp/snmpv3/view 21
name "oper"
tree " 1.3.6.1.4.1.11.2.3.7.11.33.1.2.1.3" (Agent information)
/c/sys/ssnmp/snmpv3/view 22
name "oper"
tree " 1.3.6.1.4.1.11.2.3.7.11.33.1.2.2.2" (L2 statistics)
/c/sys/ssnmp/snmpv3/view 23
name "oper"
tree " 1.3.6.1.4.1.11.2.3.7.11.33.1.2.2.3" (L2 information)
/c/sys/ssnmp/snmpv3/view 24
name "oper"
tree " 1.3.6.1.4.1.11.2.3.7.11.33.1.2.3.2" (L3 statistics)
/c/sys/ssnmp/snmpv3/view 25
name "oper"
tree " 1.3.6.1.4.1.11.2.3.7.11.33.1.2.3.3" (L3 information)

```

Configuring SNMP trap hosts

SNMPv1 trap host

1. Configure a user with no authentication or password.

```

/c/sys/ssnmp/snmpv3/usm 10 (Configure user named "vltrap")
name "vltrap"

```

2. Configure an access group and group table entries for the user. Use the following command to specify which traps can be received by the user:

```
/c/sys/ssnmp/snmpv3/access <x>/nview
```

```

/c/sys/ssnmp/snmpv3/access 10 (Define access group to view SNMPv1 traps)
name "vltrap"
model snmpv1
nview "iso"
/c/sys/ssnmp/snmpv3/group 10 (Assign user to the access group)
model snmpv1
uname vltrap
gname vltrap

```

In this example the user will receive the traps sent by the switch.

3. Configure an entry in the notify table.

```
/c/sys/ssnmp/snmpv3/notify 10    (Assign user to the notify table)
name v1trap
tag v1trap
```

4. Specify the IP address and other trap parameters in the Target Address(targetAddr) and Target Parameters (targetParam) tables. Use the following command to specify the user name used with this targetParam table:

```
c/sys/ssnmp/snmpv3/tparam <x>/uname
```

```
/c/sys/ssnmp/snmpv3/taddr 10    (Define an IP address to send traps)
name v1trap
addr 47.80.23.245
taglist v1trap
pname v1param
/c/sys/ssnmp/snmpv3/tparam 10    (Specify SNMPv1 traps to send)
name v1param
mpmodel snmpv1
uname v1trap
model snmpv1
```

5. Use the community table to define the community string used in the traps.

```
/c/sys/ssnmp/snmpv3/comm 10    (Define the community string)
index v1trap
name public
uname v1trap
```

SNMPv2 trap host configuration

The SNMPv2 trap host configuration is similar to the SNMPv1 trap host configuration. Wherever you specify the model, specify `snmpv2` instead of `snmpv1`.

```

c/sys/ssnmp/snmpv3/usm 10      (Configure user named "v2trap")
name "v2trap"
/c/sys/ssnmp/snmpv3/access 10  (Define access group to view SNMPv2 traps)
    name "v2trap"
    model snmpv2
    nview "iso"
/c/sys/ssnmp/snmpv3/group 10   (Assign user to the access group)
    model snmpv2
    uname v2trap
    gname v2trap
/c/sys/ssnmp/snmpv3/notify 10  (Assign user to the notify table)
    name v2trap
    tag v2trap
/c/sys/ssnmp/snmpv3/taddr 10   (Define an IP address to send traps)
    name v2trap
    addr 47.81.25.66
    taglist v2trap
    pname v2param
/c/sys/ssnmp/snmpv3/tparam 10  (Specify SNMPv2 traps to send)
    name v2param
    mpmodel snmpv2c
    uname v2trap
    model snmpv2
/c/sys/ssnmp/snmpv3/comm 10    (Define the community string)
    index v2trap
    name public
    uname v2trap

```

SNMPv3 trap host configuration

To configure a user for SNMPv3 traps, you can choose to send the traps with both privacy and authentication, with authentication only, or without privacy or authentication. Use the following commands to configure the access table:

```
/c/sys/ssnmp/snmpv3/access <x>/level
```

```
/c/sys/ssnmp/snmpv3/tparam <x>
```

Configure the user in the user table to match the configuration of the access table.

It is not necessary to configure the community table for SNMPv3 traps because the community string is not used by SNMPv3.

The following example shows how to configure a SNMPv3 user `v3trap` with authentication only:

```

/c/sys/ssnmp/snmpv3/usm 11      (Configure user named "v3trap")
    name "v3trap"
    auth md5
    authpw v3trap
/c/sys/ssnmp/snmpv3/access 11  (Define access group to view SNMPv3 traps)
    name "v3trap"
    level authNoPriv
    nview "iso"
/c/sys/ssnmp/snmpv3/group 11   (Assign user to the access group)
    uname v3trap
    gname v3trap
/c/sys/ssnmp/snmpv3/notify 11  (Assign user to the notify table)
    name v3trap
    tag v3trap
/c/sys/ssnmp/snmpv3/taddr 11   (Define an IP address to send traps)
    name v3trap
    addr 47.81.25.66
    taglist v3trap
    pname v3param
/c/sys/ssnmp/snmpv3/tparam 11  (Specify SNMPv3 traps to send)
    name v3param
    uname v3trap
    level authNoPriv           (Set the authentication level)

```

For more information on using SNMP, see the *HP 10Gb Ethernet BL-c Switch Command Reference Guide*.

See the *HP 10Gb Ethernet BL-c Switch User Guide* for a complete list of supported MIBs.

Secure access to the switch

Secure switch management is needed for environments that perform significant management functions across the Internet. The following are some of the functions for secured management:

- Limiting management users to a specific IP address range. See the “Setting allowable source IP address ranges” section in this chapter.
- Authentication and authorization of remote administrators. See the “RADIUS authentication and authorization” section or the “TACACS+ authentication” section, both later in this chapter.
- Encryption of management information exchanged between the remote administrator and the switch. See the “Secure Shell and Secure Copy” section later in this chapter.

Setting allowable source IP address ranges

To limit access to the switch without having to configure filters for each switch port, you can set a source IP address (or range) that will be allowed to connect to the switch IP interface through Telnet, SSH, SNMP, or the switch browser-based interface (BBI).

When an IP packet reaches the application switch, the source IP address is checked against the range of addresses defined by the management network and management mask. If the source IP address of the host or hosts is within this range, it is allowed to attempt to log in. Any packet addressed to a switch IP interface with a source IP address outside this range is discarded.

Configuring an IP address range for the management network

Configure the management network IP address and mask from the System Menu in the CLI. For example:

```
>> Main# /cfg/sys/access/mgmt/add
Enter Management Network Address: 192.192.192.0
Enter Management Network Mask: 255.255.255.128
```

In this example, the management network is set to 192.192.192.0 and management mask is set to 255.255.255.128. This defines the following range of allowed IP addresses:

192.192.192.1 to 192.192.192.127

The following source IP addresses are granted or not granted access to the switch:

- A host with a source IP address of 192.192.192.21 falls within the defined range and would be allowed to access the switch.
- A host with a source IP address of 192.192.192.192 falls outside the defined range and is not granted access. To make this source IP address valid, you would need to shift the host to an IP address within the valid range specified by the `mnet` and `mmask` or modify the `mnet` to be 192.192.192.128 and the `mmask` to be 255.255.255.128. This would put the 192.192.192.192 host within the valid range allowed by the `mnet` and `mmask` (192.192.192.128-255).

RADIUS authentication and authorization

The switch supports the Remote Authentication Dial-in User Service (RADIUS) method to authenticate and authorize remote administrators for managing the switch. This method is based on a client/server model. The Remote Access Server (RAS)—the switch—is a client to the back-end database server. A remote user (the remote administrator) interacts only with the RAS, not the back-end server and database.

RADIUS authentication consists of the following components:

- A protocol with a frame format that utilizes User Datagram Protocol (UDP) over IP, based on Request For Comments (RFC) 2138 and 2866
- A centralized server that stores all the user authorization information
- A client, in this case, the switch

The switch, acting as the RADIUS client, communicates to the RADIUS server to authenticate and authorize a remote administrator using the protocol definitions specified in RFC 2138 and 2866. Transactions between the client and the RADIUS server are authenticated using a shared key that is not sent over the network. In addition, the remote administrator passwords are sent encrypted between the RADIUS client (the switch) and the back-end RADIUS server.

How RADIUS authentication works

RADIUS authentication works as follows:

1. A remote administrator connects to the switch and provides the user name and password.
2. Using Authentication/Authorization protocol, the switch sends the request to the authentication server.
3. The authentication server checks the request against the user ID database.
4. Using RADIUS protocol, the authentication server instructs the switch to grant or deny administrative access.

Configuring RADIUS on the switch (CLI example)

To configure RADIUS on the switch, do the following:

1. Turn RADIUS authentication on, and then configure the Primary and Secondary RADIUS servers. For example:

```
>> Main# /cfg/sys/radius          (Select the RADIUS Server menu)
>> RADIUS Server# on              (Turn RADIUS on)
Current status: OFF
New status:      ON
>> RADIUS Server# prisrv 10.10.1.1 (Enter primary server IP)
Current primary RADIUS server:    0.0.0.0
New pending primary RADIUS server: 10.10.1.1
>> RADIUS Server# secsrv 10.10.1.2 (Enter secondary server IP)
Current secondary RADIUS server:   0.0.0.0
New pending secondary RADIUS server: 10.10.1.2
```

2. Configure the primary RADIUS secret and secondary RADIUS secret.

```
>> RADIUS Server# secret
Enter new RADIUS secret: <1-32 character secret>
>> RADIUS Server# secret2
Enter new RADIUS second secret: <1-32 character secret>
```



CAUTION: If you configure the RADIUS secret using any method other than a direct console connection, the secret may be transmitted over the network as clear text.

3. If desired, you may change the default User Datagram Protocol (UDP) port number used to listen to RADIUS. The well-known port for RADIUS is 1645.

```
>> RADIUS Server# port
Current RADIUS port: 1645
Enter new RADIUS port [1500-3000]: <UDP port number>
```

4. Configure the number of retry attempts for contacting the RADIUS server and the timeout period.


```
>> RADIUS Server# retries
Current RADIUS server retries: 3
Enter new RADIUS server retries [1-3]:<server retries>
>> RADIUS Server# time
Current RADIUS server timeout: 3
Enter new RADIUS server timeout [1-10]: 10 (Enter the timeout period
in seconds)
```

5. Apply and save the configuration.

```
>> RADIUS Server# apply
>> RADIUS Server# save
```

Configuring RADIUS on the switch (BBI example)

1. Configure RADIUS parameters.
 - a. Click the **Configure** context button.
 - b. Open the System folder, and select **Radius**.



- c. Enter the IP address of the primary and secondary RADIUS servers, and enter the RADIUS secret for each server. Enable the RADIUS server.

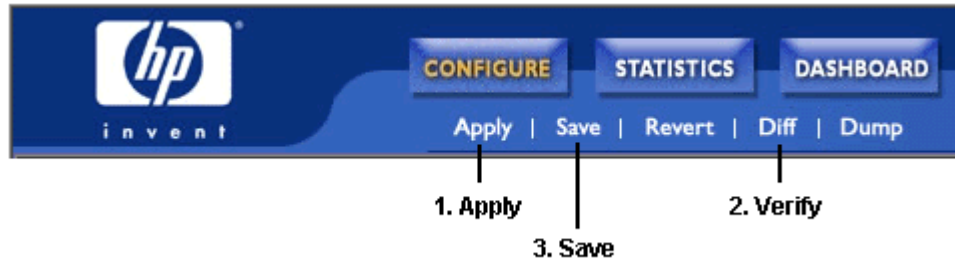
Switch Radius Configuration

Primary Radius IP Address	<input type="text" value="10.10.2.1"/>
Secondary Radius IP Address	<input type="text" value="10.10.2.2"/>
Radius port (1500-3000)	<input type="text" value="1645"/>
Radius timeout (1-10)	<input type="text" value="3"/>
Radius retries (1-3)	<input type="text" value="3"/>
Enable/Disable Radius Server	<input type="text" value="Enabled"/> ▼
Enable/Disable Radius Backdoor for telnet/ssh/http/https	<input type="text" value="Disabled"/> ▼
Enable/Disable Radius Secure Backdoor for telnet/ssh/http/https	<input type="text" value="Disabled"/> ▼
Radius Secret	<input type="text"/>
Secondary Radius Server Secret	<input type="text"/>

CAUTION: If you configure the RADIUS secret using any method other than a direct console connection, the secret may be transmitted over the network as clear text.

- d. Click **Submit**.

2. Apply, verify, and save the configuration.



RADIUS authentication features

The switch supports the following RADIUS authentication features:

- Supports RADIUS client on the switch, based on the protocol definitions in RFC 2138 and RFC 2866.
- Allows RADIUS secret password up to 32 bytes.
- Supports secondary authentication server so that when the primary authentication server is unreachable, the switch can send client authentication requests to the secondary authentication server. Use the `/cfg/sys/radius/cur` command to show the currently active RADIUS authentication server.
- Supports user-configurable RADIUS server retry and time-out values:
 - Time-out value = 1-10 seconds
 - Retries = 1-3
- The switch will time out if it does not receive a response from the RADIUS server in one to three retries. The switch will also automatically retry connecting to the RADIUS server before it declares the server down.
- Supports user-configurable RADIUS application port. The default is 1645/User Datagram Protocol (UDP)-based on RFC 2138. Port 1812 is also supported.
- Allows network administrator to define privileges for one or more specific users to access the switch at the RADIUS user database.
- Allows the administrator to configure RADIUS backdoor and secure backdoor for Telnet, SSH, HTTP, and HTTPS access.

User accounts for RADIUS users

The user accounts listed in the following table can be defined in the RADIUS server dictionary file.

Table 2 User access levels

User account	Description and tasks performed
User	User interaction with the switch is completely passive; nothing can be changed on the switch. Users may display information that has no security or privacy implications, such as switch statistics and current operational state information.
Operator	Operators can only effect temporary changes on the switch. These changes are lost when the switch is rebooted/reset. Operators have access to the switch management features used for daily switch operations. Because any changes an operator makes are undone by a reset of the switch, operators cannot severely impact switch operation, but do have access to the Maintenance menu. By default, the operator account is disabled and has no password.

Table 2 User access levels

User account	Description and tasks performed
Administrator	Administrators are the only ones that can make permanent changes to the switch configuration—changes that are persistent across a reboot/reset of the switch. Administrators can access switch functions to configure and troubleshoot problems on the switch level. Because administrators can also make temporary (operator-level) changes as well, they must be aware of the interactions between temporary and permanent changes.

RADIUS attributes for user privileges

When the user logs in, the switch authenticates the level of access by sending the RADIUS access request, that is, the client authentication request, to the RADIUS authentication server.

If the authentication server successfully authenticates the remote user, the switch verifies the privileges of the remote user and authorizes the appropriate access. The administrator has the option to allow backdoor access through the console port only, or through the console and Telnet/SSH/HTTP/HTTPS access. When backdoor access is enabled, access is allowed even if the primary and secondary authentication servers are reachable. Only when both the primary and secondary authentication servers are not reachable, the administrator has the option to allow secure backdoor (`secbã`) access through the console port only, or through the console and Telnet/SSH/HTTP/HTTPS access. When RADIUS is on, you can have either backdoor or secure backdoor enabled, but not both at the same time. The default value for backdoor access through the console port only is `enabled`. You always can access the switch via the console port, by using `noradius` and the administrator password, whether backdoor/secure backdoor are enabled or not. The default value for backdoor and secure backdoor access through Telnet/SSH/HTTP/HTTPS is `disabled`.

All user privileges, other than those assigned to the administrator, must be defined in the RADIUS dictionary. RADIUS attribute 6, which is built into all RADIUS servers, defines the administrator. The file name of the dictionary is RADIUS vendor-dependent. The RADIUS attributes shown in the following table are defined for user privilege levels.

Table 3 Proprietary attributes for RADIUS

User name/access	User service type	Value
User	Vendor-supplied	255
Operator	Vendor-supplied	252

TACACS+ authentication

The switch software supports authentication, authorization, and accounting with networks using the Cisco Systems TACACS+ protocol. The switch functions as the Network Access Server (NAS) by interacting with the remote client and initiating authentication and authorization sessions with the TACACS+ access server. The remote user is defined as someone requiring management access to the switch either through a data or management port.

TACACS+ offers the following advantages over RADIUS:

- TACACS+ uses TCP-based connection-oriented transport; whereas RADIUS is UDP based. TCP offers a connection-oriented transport, while UDP offers best-effort delivery. RADIUS requires additional programmable variables such as re-transmit attempts and time-outs to compensate for best-effort transport, but it lacks the level of built-in support that a TCP transport offers.
- TACACS+ offers full packet encryption whereas RADIUS offers password-only encryption in authentication requests.
- TACACS+ separates authentication, authorization, and accounting.

How TACACS+ authentication works

TACACS+ works much in the same way as RADIUS authentication.

1. Remote administrator connects to the switch and provides user name and password.

NOTE: The user name and password can have a maximum length of 128 characters. The password cannot be left blank.

2. Using Authentication/Authorization protocol, the switch sends request to authentication server.
3. Authentication server checks the request against the user ID database.
4. Using TACACS+ protocol, the authentication server instructs the switch to grant or deny administrative access.

During a session, if additional authorization checking is needed, the switch checks with a TACACS+ server to determine if the user is granted permission to use a particular command.

TACACS+ authentication features

Authentication is the action of determining the identity of a user, and is generally done when the user first attempts to log in to a device or gain access to its services. Switch software supports ASCII inbound login to the device. PAP, CHAP and ARAP login methods, TACACS+ change password requests, and one-time password authentication are not supported.

Authorization

Authorization is the action of determining a user's privileges on the device, and usually takes place after authentication.

The default mapping between TACACS+ authorization privilege levels and switch management access levels is shown in the table below. The privilege levels listed in the following table must be defined on the TACACS+ server.

Table 4 Default TACACS+ privilege levels

User access level	TACACS+ level
user	0
oper	3
admin	6

Alternate mapping between TACACS+ privilege levels and HP 10GbE switch management access levels is shown in the table below. Use the command `/cfg/sys/tacacs/cmap ena` to use the alternate TACACS+ privilege levels.

Table 5 Alternate TACACS+ privilege levels

User access level	TACACS+ level
user	0–1
oper	6–8
admin	14–15

You can customize the mapping between TACACS+ privilege levels and HP 10GbE switch management access levels. Use the command `/cfg/sys/tacacs/usermap` to manually map each TACACS+ privilege level (0-15) to a corresponding HP 10GbE switch management access level (user, oper, admin, none).

If the remote user is authenticated by the authentication server, the HP 10GbE switch verifies the privileges of the remote user and authorizes the appropriate access. When both the primary and secondary authentication servers are not reachable, the administrator has an option to allow backdoor access via the console only or console and Telnet access. The default value is `disable` for Telnet access and `enable` for console access. The administrator also can enable secure backdoor (`/cfg/sys/tacacs/secbd`) to allow access if both the primary and secondary TACACS+ servers fail to respond.

Accounting

Accounting is the action of recording a user's activities on the device for the purposes of billing and/or security. It follows the authentication and authorization actions. If the authentication and authorization is not performed via TACACS+, no TACACS+ accounting messages are sent out.

You can use TACACS+ to record and track software logins, configuration changes, and interactive commands.

The switch supports the following TACACS+ accounting attributes:

- protocol (console/telnet/ssh/http)
- start_time
- stop_time
- elapsed_time

NOTE: When using the browser-based Interface, the TACACS+ Accounting Stop records are sent only if the **Quit** button on the browser is clicked.

Configuring TACACS+ authentication on the switch (CLI example)

1. Turn TACACS+ authentication on, and then configure the Primary and Secondary TACACS+ servers.

```
>> Main# /cfg/sys/tacacs          (Select the TACACS+ Server menu)
>> TACACS+ Server# on             (Turn TACACS+ on)
Current status: OFF
New status: ON
>> TACACS+ Server# prisrv 10.10.1.1 (Enter primary server IP)
Current primary TACACS+ server: 0.0.0.0
New pending primary TACACS+ server: 10.10.1.1
>> TACACS+ Server# secsrv 10.10.1.2 (Enter secondary server IP)
Current secondary TACACS+ server: 0.0.0.0
New pending secondary TACACS+ server: 10.10.1.2
```

2. Configure the TACACS+ secret and second secret.

```
>> TACACS+ Server# secret
Enter new TACACS+ secret: <1-32 character secret>
>> TACACS+ Server# secret2
Enter new TACACS+ second secret: <1-32 character secret>
```



CAUTION: If you configure the TACACS+ secret using any method other than a direct console connection, the secret may be transmitted over the network as clear text.

3. If desired, you may change the default TCP port number used to listen to TACACS+. The well-known port for TACACS+ is 49.

```
>> TACACS+ Server# port
Current TACACS+ port: 49
Enter new TACACS+ port [1-65000]: <TCP port number>
```

4. Configure the number retry attempts for contacting the TACACS+ server and the timeout period.

```
>> TACACS+ Server# retries
Current TACACS+ server retries: 3
Enter new TACACS+ server retries [1-3]: 2
>> TACACS+ Server# time
Current TACACS+ server timeout: 5
Enter new TACACS+ server timeout [4-15]: 10 (Enter the timeout period
                                             in minutes)
```

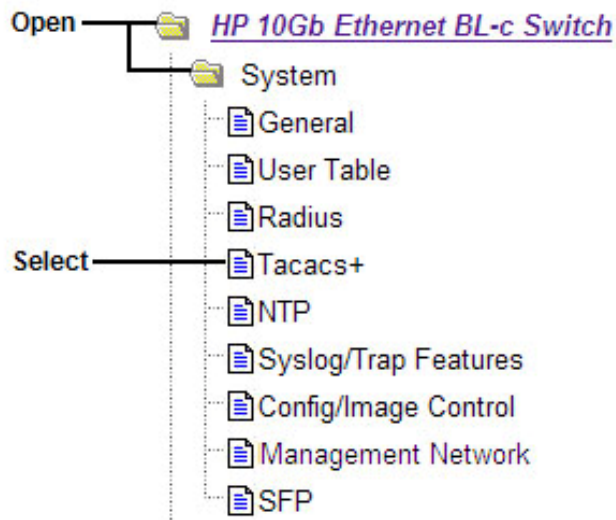
5. Configure custom privilege-level mapping (optional).

```
>> TACACS+ Server# usermap 2
Current privilege mapping for remote privilege 2: not set
Enter new local privilege mapping: user
>> TACACS+ Server# usermap 3 user
>> TACACS+ Server# usermap 4 user
>> TACACS+ Server# usermap 5 oper
```

6. Apply and save the configuration.

Configuring TACACS+ authentication on the switch (BBI example)

1. Configure TACACS+ authentication for the switch.
 - a. Click the **Configure** context button.
 - b. Open the System folder, and select **Tacacs+**.



- c. Enter the IP address of the primary and secondary TACACS+ servers, and enter the TACACS+ secret. Enable TACACS+.

Switch Tacacs+ Configuration	
Primary Tacacs+ IP Address	<input type="text" value="10.10.1.1"/>
Secondary Tacacs+ IP Address	<input type="text" value="10.10.1.2"/>
Tacacs+ port (1-65000)	<input type="text" value="49"/>
Tacacs+ timeout (4-15)	<input type="text" value="5"/>
Tacacs+ retries (1-3)	<input type="text" value="3"/>
Enable/Disable Tacacs+ Server	Enabled <input type="button" value="v"/>
Enable/Disable Tacacs+ Backdoor for telnet/ssh/http/https	Disabled <input type="button" value="v"/>
Enable/Disable Tacacs+ Secure Backdoor for telnet/ssh/http/https	Disabled <input type="button" value="v"/>
Enable/Disable Tacacs+ new privilege level mapping	Disabled <input type="button" value="v"/>
Tacacs+ Secret	<input type="text"/>
Secondary Tacacs+ Server Secret	<input type="text"/>

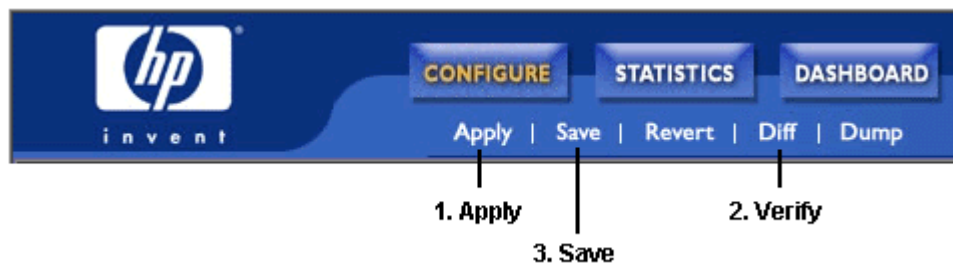
- d. Click **Submit**.

- e. Configure custom privilege-level mapping (optional). Click **Submit** to accept each mapping change.

Tacacs+ User Mappings Configuration

Remote privilege	Local privilege
5	Oper
0	not set
1	user
2	user
3	user
4	user
5	not set
⋮	
14	not set
15	not set

2. Apply, verify, and save the configuration.



Secure Shell and Secure Copy

Secure Shell (SSH) and Secure Copy (SCP) use secure tunnels to encrypt and secure messages between a remote administrator and the switch. Telnet does not provide this level of security. The Telnet method of managing a switch does not provide a secure connection.

SSH is a protocol that enables remote administrators to log securely into the switch over a network to execute management commands. By default, SSH is disabled (off) on the switch.

SCP is typically used to copy files securely from one machine to another. SCP uses SSH for encryption of data on the network. On a switch, SCP is used to download and upload the switch configuration via secure channels. By default, SCP is disabled on the switch.

The switch implementation of SSH is based on version 1.5 and version 2.0, and supports SSH clients from version 1.0 through version 2.0. Client software can use SSH version 1 or version 2. The following SSH clients are supported:

- SSH 3.0.1 for Linux (freeware)
- SecureCRT® 4.1.8 (VanDyke Technologies, Inc.)
- OpenSSH_3.9 for Linux (FC 3)
- FedoraCore 3 for SCP commands
- PuTTY Release 0.58 (Simon Tatham) for Windows

Configuring SSH and SCP features (CLI example)

Before you can use SSH commands, use the following commands to turn on SSH and SCP.

Enabling or disabling SSH

To enable the SSH feature, connect to the switch CLI and enter the following commands:

```
>> # /cfg/sys/sshd/on                (Turn SSH on)
Current status: OFF
New status: ON
SSHD# apply                          (Apply the changes to start generating
                                      RSA host and server keys)

RSA host key generation starts
. . . . .
. . . . .
RSA host key generation completes (lasts 212549 ms)
RSA host key is being saved to Flash ROM, please don't reboot the box
immediately.
RSA server key generation starts
. . . . .
. . . . .
RSA server key generation completes (lasts 75503 ms)
RSA server key is being saved to Flash ROM, please don't reboot the box
immediately.
-----
Apply complete; don't forget to "save" updated configuration.
```

Enabling or disabling SCP apply and save

Enter the following commands from the switch CLI to enable the SCP `putcfg_apply` and `putcfg_apply_save` commands:

```
>> # /cfg/sys/sshd/ena                (Enable SCP apply and save)
>> # /cfg/sys/sshd/dis                (Disable SCP apply and save)
SSHD# apply                          (Apply the changes)
```

Configuring the SCP administrator password

To configure the `scpadmin` (SCP administrator) password, first connect to the switch via the RS-232 management console. For security reasons, the `scpadmin` password can be configured only when connected directly to the switch console.

To configure the password, enter the following CLI command. At factory default settings, the current SCP administrator password is `admin`.

```
>> # /cfg/sys/sshd/scpadm
Changing SCP-only Administrator password; validation required. . .
Enter current administrator password: <password>
Enter new SCP-only administrator password: <new password>
Re-enter new SCP-only administrator password: <new password>
New SCP-only administrator password accepted.
```



IMPORTANT: The SCP-only administrator password must be different from the regular administrator password.

Using SSH and SCP client commands

The following shows the format for using some client commands. The examples below use 205.178.15.157 as the IP address of a sample switch.

Logging in to the switch

Enter the following command to log in to the switch:

```
ssh <user>@<switch IP address>
```

For example:

```
>> # ssh admin@205.178.15.157
```

Downloading configuration from the switch using SCP

Enter the following command to download the switch configuration using SCP. You will be prompted for a password:

```
scp <user>@<switch IP address>:getcfg <local filename>
```

For example:

```
>> # scp scpadmin@205.178.15.157:getcfg ad4.cfg
```

The switch prompts you for the `scpadmin` password.

Uploading configuration to the switch using SCP

Enter the following command to upload configuration to the switch. You will be prompted for a password.

```
scp <local filename> <user>@<switch IP address>:putcfg
```

For example:

```
>> # scp ad4.cfg admin@205.178.15.157:putcfg
```


Applying and saving configuration

Enter the apply and save commands after the command above (`scp ad4.cfg 205.178.15.157:putcfg`), or use the following commands. You will be prompted for a password.

```
>> # scp <local_filename> <user>@<switch IP addr>:putcfg_apply
>> # scp <local_filename> <user>@<switch IP addr>:putcfg_apply_save
```

For example:

```
>> # scp ad4.cfg admin@205.178.15.157:putcfg_apply
>> # scp ad4.cfg admin@205.178.15.157:putcfg_apply_save
```

NOTE:

- The `diff` command is automatically executed at the end of `putcfg` to notify the remote client of the difference between the new and the current configurations.
 - `putcfg_apply` runs the apply command after the `putcfg` is done.
 - `putcfg_apply_save` saves the new configuration to the flash after `putcfg_apply` is done.
 - The `putcfg_apply` and `putcfg_apply_save` commands are provided because extra apply and save commands are usually required after a `putcfg`.
-

SSH and SCP encryption of management messages

The following encryption and authentication methods are supported for SSH and SCP:

- Server Host Authentication—Client RSA authenticates the switch at the beginning of every connection
- Key Exchange—RSA
- Encryption—AES256-CBC, AES192-CBC, 3DES-CBC, 3DES, ARCFOUR
- User Authentication—Local password authentication, RADIUS, TACACS+

Generating RSA host and server keys for SSH access

To support the SSH server feature, two sets of RSA keys (host and server keys) are required. The host key is 1024 bits and is used to identify the switch. The server key is 768 bits and is used to make it impossible to decipher a captured session by breaking into the switch at a later time.

When the SSH server is first enabled and applied, the switch automatically generates the RSA host and server keys and is stored in the flash memory.

To configure RSA host and server keys, first connect to the switch console connection (commands are not available via Telnet connection), and enter the following commands to generate them manually:

```
>> # /cfg/sys/sshd/hkeygen (Generates the host key)
>> # /cfg/sys/sshd/skeygen (Generates the server key)
```

These two commands take effect immediately without the need of an `apply` command.

When the switch reboots, it will retrieve the host and server keys from the flash memory. If these two keys are not available in the flash memory and if the SSH server feature is enabled, the switch automatically generates them during the system reboot. This process may take several minutes to complete.

The switch can also automatically regenerate the RSA server key. To set the interval of RSA server key auto-generation, use the following command:

```
>> # /cfg/sys/sshd/intrval <number of hours (0-24)>
```

A value of 0 denotes that RSA server key autogeneration is disabled. When greater than 0, the switch will auto generate the RSA server key every specified interval; however, RSA server key generation is skipped if the switch is busy doing other key or cipher generation when the timer expires.

The switch will perform only one session of key/cipher generation at a time. Thus, an SSH/SCP client will not be able to log in if the switch is performing key generation at that time, or if another client has logged in immediately prior. Also, key generation will fail if an SSH/SCP client is logging in at that time.

SSH/SCP integration with RADIUS and TACACS+ authentication

SSH/SCP is integrated with RADIUS and TACACS+ authentication. After the RADIUS or TACACS+ server is enabled on the switch, all subsequent SSH authentication requests will be redirected to the specified RADIUS or TACACS+ servers for authentication. The redirection is transparent to the SSH clients.

User access control

The switch allows an administrator to define end user accounts that permit end users to perform limited actions on the switch. Once end user accounts are configured and enabled, the switch requires username/password authentication.

For example, an administrator can assign a user who can log into the switch and perform operational commands (effective only until the next switch reboot).

The administrator defines access levels for each switch user, as shown in the following table.

Table 6 User access levels

User account	Description	Password
Administrator	The Administrator has complete access to all menus, information, and configuration commands on the switch, including the ability to change both the user and administrator passwords.	admin
Operator	The Operator manages all functions of the switch. The Operator can reset ports or the entire switch.	oper
User	The User has no direct responsibility for switch management. Users can view all switch status information and statistics but cannot make any configuration changes to the switch.	user

Passwords can be up to 128 characters in length for TACACS+, Telnet, SSH, console, and BBI access. When RADIUS authentication is used, the maximum password length is 32 characters.

If RADIUS authentication is used, the user password on the Radius server will override the user password on the switch. Also note that the password-change command on the switch modifies only the `user switch` password and has no effect on the user password on the Radius server. RADIUS authentication and user password cannot be used concurrently to access the switch.

Setting up user IDs

The administrator can configure up to 10 user accounts.

To configure an end-user account, perform the following steps:

1. Select a user ID to define.

```
>> # /cfg/sys/access/user/uid 1
```

2. Define the user name and password.

```
>> User ID 1 # name jane (Assign name "jane" to user ID 1)
Current user name:
New user name: jane
```

3. Define the user access level. By default, the end user is assigned to the user access level. To change the user's access level, enter the user Class of Service (CoS) command, and select one of the available options.

```
>> User ID 1 # cos <user|oper|admin>
```

4. Enable the user ID.

```
>> # /cfg/sys/access/user/uid <#>/ena
```

Once an end user account is configured and enabled, the user can login to the switch using the username/password combination. The level of switch access is determined by the user CoS for the account. The CoS corresponds to the user access levels described in the User access levels table.

Ports and trunking

Introduction

The first part of this chapter describes the different types of ports used on the switch. This information is useful in understanding other applications described in this guide, from the context of the embedded switch/server environment.

For specific information on how to configure ports for speed, auto-negotiation, and duplex modes, see the port commands in the *HP 10Gb Ethernet BL-c Switch Command Reference Guide*.

The second part of this chapter provides configuration background and examples for trunking multiple ports together. Trunk groups can provide super-bandwidth, multi-link connections between switches or other trunk-capable devices. A trunk group is a group of links that act together, combining their bandwidth to create a single, larger virtual link. The switch provides trunking support for the four external ports and 16 server ports.

Ports on the switch

The following table describes the Ethernet ports of the switch, including the port name and function.

NOTE: The actual mapping of switch ports to NIC interfaces is dependant on the operating system software, the type of server blade, and the enclosure type. For more information, see the *HP 10Gb Ethernet BL-c Switch User Guide*.

Table 7 Ethernet switch port names

Port number	Port alias
1	Downlink1
2	Downlink2
3	Downlink3
4	Downlink4
5	Downlink5
6	Downlink6
7	Downlink7
8	Downlink8
9	Downlink9
10	Downlink10
11	Downlink11
12	Downlink12

Table 7 Ethernet switch port names

Port number	Port alias
13	Downlink13
14	Downlink14
15	Downlink15
16	Downlink16
17	Mgmt
18	Uplink1
19	Uplink2
20	Uplink3
21	Uplink4

Port trunk groups

When using port trunk groups between two switches, you can create an aggregate link operating at up to forty Gigabits per second, depending on how many physical ports are combined. The switch supports up to 12 trunk groups per switch, each with up to six ports per trunk group.

The trunking software detects broken trunk links (link down or disabled) and redirects traffic to other trunk members within that trunk group. You can only use trunking if each link has the same configuration for speed, flow control, and auto-negotiation.

Statistical load distribution

In a configured trunk group containing more than one port, the load distribution is determined by information embedded within the data frame. For IP traffic, the switch will calculate the trunk port to use for forwarding traffic by implementing the load distribution algorithm on value equals to modulus of (XOR of last 3 bits of Source and last 3 bits of Destination IP address). For non-IP traffic, the switch will calculate the trunk port to use for forwarding traffic by implementing the load distribution algorithm on value equals to modulus of (XOR of last 3 bits of Source and last 3 bits of Destination MAC address).

Built-in fault tolerance

Since each trunk group is composed of multiple physical links, the trunk group is inherently fault tolerant. As long as even one physical link between the switches is available, the trunk remains active.

Statistical load distribution is maintained whenever a link in a trunk group is lost or returned to service.

Before you configure trunks

When you create and enable a trunk, the trunk members (switch ports) take on certain settings necessary for correct operation of the trunking feature.

Before you configure your trunk, you must consider these settings, along with specific configuration rules, as follows:

1. Read the configuration rules provided in the “Trunk group configuration rules” section.
2. Determine which switch ports (up to six) are to become trunk members (the specific ports making up the trunk).
3. Ensure that the chosen switch ports are set to enabled, using the following command:
`/cfg/port x/cur`
4. Trunk member ports must have the same VLAN configuration.
5. Consider how the existing spanning tree will react to the new trunk configuration. See the “Spanning Tree Protocol” chapter for spanning tree group configuration guidelines.
6. Consider how existing VLANs will be affected by the addition of a trunk.

Trunk group configuration rules

The trunking feature operates according to specific configuration rules. When creating trunks, consider the following rules that determine how a trunk group reacts in any network topology:


- All trunks must originate from one device, and lead to one destination device. For example, you cannot combine a link from Server 1 and a link from Server 2 into one trunk group.
- Any physical switch port can belong to only one trunk group.
- Trunking from non-HP devices must comply with Cisco® EtherChannel® technology.
- All trunk member ports must be assigned to the same VLAN configuration before the trunk can be enabled.
- If you change the VLAN settings of any trunk member, you cannot apply the change until you change the VLAN settings of all trunk members.
- When an active port is configured in a trunk, the port becomes a trunk member when you enable the trunk using the `/cfg/l2/trunk x/ena` command. The spanning tree parameters for the port then change to reflect the new trunk settings.
- All trunk members must be in the same spanning tree group and can belong to only one spanning tree group. However if all ports are tagged, then all trunk ports can belong to multiple spanning tree groups.
- When a trunk is enabled, the trunk spanning tree participation setting takes precedence over that of any trunk member.
- You cannot configure a trunk member as a monitor port in a Port Mirroring configuration.
- A monitor port cannot monitor trunks; however, trunk members can be monitored.

Port trunking example

In this example, the 10 Gigabit uplink ports on each switch are configured into a total of four trunk groups: two on each switch.

NOTE: The actual mapping of switch ports to NIC interfaces is dependant on the operating system software, the type of server blade, and the enclosure type. For more information, see the *HP 10Gb Ethernet BL-c Switch User Guide*.

Figure 1 Port trunk group configuration example



The trunk groups are configured as follows:

Trunk groups 2-5 consist of two 10 Gigabit uplink ports each, configured to act as a single link to the upstream routers. The trunk groups on each switch are configured so that there is a link to each router for redundancy.

Prior to configuring each switch in this example, you must connect to the appropriate switch CLI as the administrator. For details about accessing and using any of the commands described in this example, see the *HP 10Gb Ethernet BL-c Switch Command Reference*.

Configuring trunk groups (CLI example)

1. On Switch 1, configure trunk groups 5 and 3:

```
>> # /cfg/l2/trunk 5                (Select trunk group 5)
>> Trunk group 5# add 20             (Add port 20 to trunk group 5)
>> Trunk group 5# add 21             (Add port 21 to trunk group 5)
>> Trunk group 5# ena                (Enable trunk group 5)
>> Trunk group 5# apply              (Make your changes active)

>> # /cfg/l2/trunk 3                (Select trunk group 3)
>> Trunk group 3# add 18             (Add port 18 to trunk group 3)
>> Trunk group 3# add 19             (Add port 19 to trunk group 3)
>> Trunk group 3# ena                (Enable trunk group 3)
>> Trunk group 3# apply              (Make your changes active)
>> Trunk group 3# save                (Save for restore after reboot)
```

2. On Switch 2, configure trunk groups 4 and 2:

```
>> # /cfg/l2/trunk 4                (Select trunk group 4)
>> Trunk group 4# add 20             (Add port 20 to trunk group 4)
>> Trunk group 4# add 21             (Add port 21 to trunk group 4)
>> Trunk group 4# ena                (Enable trunk group 4)
>> Trunk group 4# apply              (Make your changes active)

>> # /cfg/l2/trunk 2                (Select trunk group 2)
>> Trunk group 2# add 18             (Add port 18 to trunk group 2)
>> Trunk group 2# add 19             (Add port 19 to trunk group 2)
>> Trunk group 2# ena                (Enable trunk group 2)
>> Trunk group 2# apply              (Make your changes active)
>> Trunk group 2# save                (Save for restore after reboot)
```

NOTE: In this example, two switches are used. Any third-party device supporting link aggregation should be configured manually. Connection problems might arise when using automatic trunk group negotiation on the third-party device.

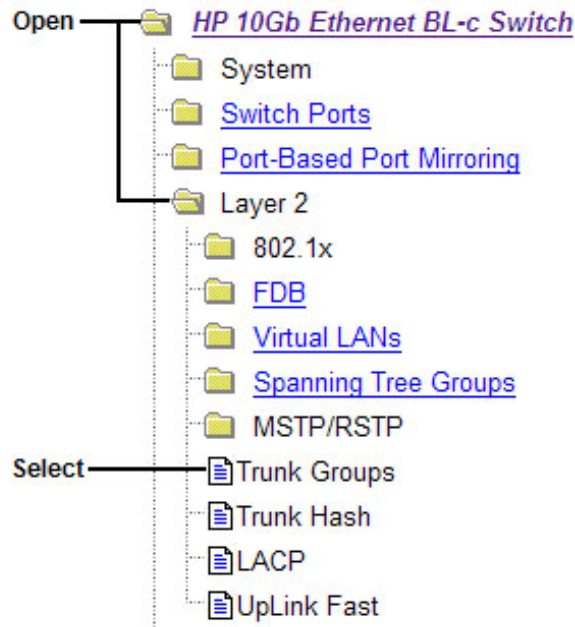
3. Examine the trunking information on each switch, using the following command:

```
>> /info/l2/trunk                    (View trunking information)
```

Information about each port in each configured trunk group is displayed. Make sure that trunk groups consist of the expected ports and that each port is in the expected state.

Configuring trunk groups (BBI example)

1. Configure trunk groups.
 - a. Click the **Configure** context button on the Toolbar.
 - b. Open the Layer 2 folder, and select **Trunk Groups**.



- c. Click a Trunk Group number to select it.

Trunk Groups Configuration

Trunk Group	State
<u>1</u>	enabled
<u>2</u>	disabled
<u>3</u>	disabled
<u>4</u>	disabled
<u>5</u>	disabled
<u>6</u>	disabled
<u>7</u>	disabled
<u>8</u>	disabled
<u>9</u>	disabled
<u>10</u>	disabled
<u>11</u>	disabled
<u>12</u>	disabled

Select

- d. Enable the Trunk Group. To add ports, select each port in the Ports Available list, and click **Add**.

Switch Trunk Group 5 Configuration

Trunk State: Enabled

Ports Available

Port

- 1
- 2
- 3
- 4

Ports added to Trunk

Port

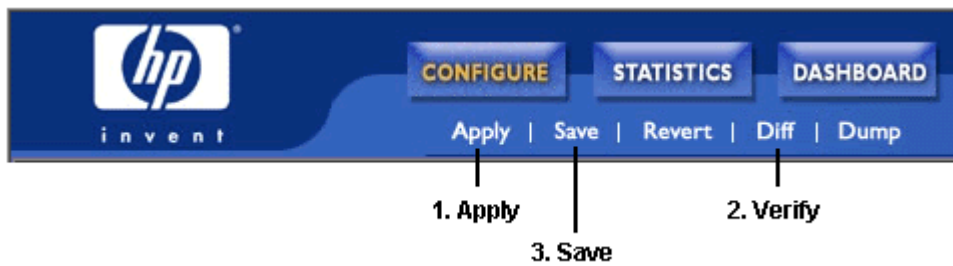
- 20
- 21

Add>>
<<Remove

Submit
Delete


- e. Click **Submit**.

2. Apply, verify, and save the configuration.

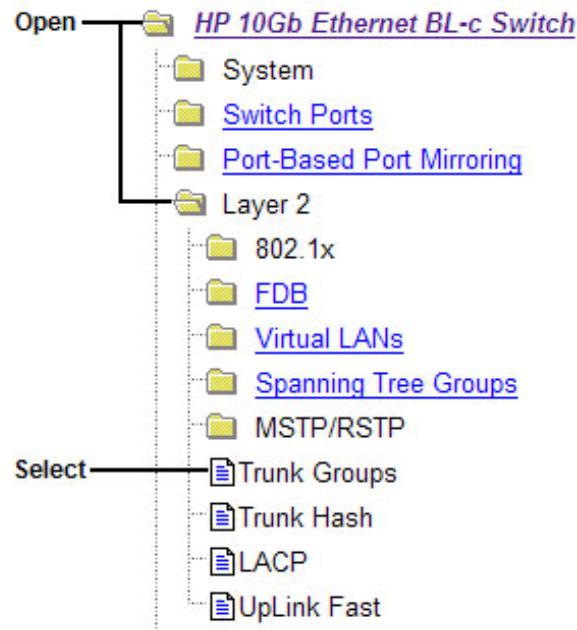


3. Examine the trunking information on each switch.

- a. Click the **Dashboard** context button on the Toolbar.



- b. Select Trunk Groups.



- c. Information about each configured trunk group is displayed. Make sure that trunk groups consist of the expected ports and that each port is in the expected state.

Switch Trunk Groups Dashboard

Status	Trunk Group	Switch Port	STG
	3 status: enabled	18	1
	3 status: enabled	19	1
	5 status: enabled	20	1
	5 status: enabled	21	1

Legend Info

	Port is down
	Port is in forwarding state
	Port is in blocking state

Configurable Trunk Hash algorithm

This feature allows you to configure the particular parameters for the HP 10GbE switch Trunk Hash algorithm instead of having to utilize the defaults. You can configure new default behavior for Layer 2 traffic and Layer 3 traffic, using the CLI menu `cfg/12/trash`. You can select a minimum of one or a maximum of two parameters to create one of the following configurations:

- Source IP (SIP)
- Destination IP (DIP)
- Source MAC (SMAC)
- Destination MAC (DMAC)
- Source IP (SIP) + Destination IP (DIP)
- Source MAC (SMAC) + Destination MAC (DMAC)

Link Aggregation Control Protocol

Link Aggregation Control Protocol (LACP) is an IEEE 802.3ad standard for grouping several physical ports into one logical port (known as a dynamic trunk group or Link Aggregation group) with any device that supports the standard. Refer to the IEEE 802.3ad-2002 for a full description of the standard.

The 802.3ad standard allows standard Ethernet links to form a single Layer 2 link using the Link Aggregation Control Protocol (LACP). Link aggregation is a method of grouping physical link segments of the same media type and speed in full duplex, and treating them as if they were part of a single, logical link segment. If a link in a LACP trunk group fails, traffic is reassigned dynamically to the remaining link(s) of the dynamic trunk group.

NOTE: Currently, LACP implementation does not support the Churn machine, an option used to detect if the port is operable within a bounded time period between the actor and the partner. Only the Market Responder is implemented, and there is no marker protocol generator.

A port's Link Aggregation Identifier (LAG ID) determines how the port can be aggregated. The Link Aggregation ID (LAG ID) is constructed mainly from the system ID and the port's admin key, as follows:

- System ID—The system ID is an integer value based on the switch's MAC address and the system priority assigned in the CLI.
- Admin key—A port's admin key is an integer value (1-65535) that you can configure in the CLI. Each HP 10GbE switch port that participates in the same LACP trunk group must have the same admin key value. The admin key is local significant, which means the partner switch does not need to use the same admin key value.

For example, consider two switches, an Actor (the HP 10GbE switch) and a Partner (another switch), as shown in the following table:

Table 8 Actor vs. partner LACP configuration

Actor Switch	Partner Switch 1	Partner Switch 2
Port 18 (admin key = 100)	Port 1 (admin key = 50)	
Port 19 (admin key = 100)	Port 2 (admin key = 50)	
Port 20 (admin key = 200)		Port 3 (admin key = 60)
Port 21 (admin key = 200)		Port 4 (admin key = 60)

In the configuration shown in the table above, Actor switch ports 18 and 19 aggregate to form an LACP trunk group with Partner switch ports 1 and 2. At the same time, Actor switch ports 20 and 21 form a different LACP trunk group with a different partner.

LACP automatically determines which member links can be aggregated and then aggregates them. It provides for the controlled addition and removal of physical links for the link aggregation.

Each port in the HP 10GbE switch can have one of the following LACP modes.

- off (default)—The user can configure this port in to a regular static trunk group.
- active—The port is capable of forming an LACP trunk. This port sends LACPDU packets to partner system ports.
- passive—The port is capable of forming an LACP trunk. This port only responds to the LACPDU packets sent from an LACP active port.

Each active LACP port transmits LACP data units (LACPDUs), while each passive LACP port listens for LACPDUs. During LACP negotiation, the admin key is exchanged. The LACP trunk group is enabled as long as the information matches at both ends of the link. If the admin key value changes for a port at either end of the link, that port's association with the LACP trunk group is lost.

When the system is initialized, all ports by default are in LACP `off` mode and are assigned unique admin keys. To make a group of ports aggregatable, you assign them all the same admin key. You must set the port's LACP mode to active to activate LACP negotiation. You can set other port's LACP mode to passive, to reduce the amount of LACPDU traffic at the initial trunk-forming stage.

Use the `/info/12/trunk` command or the `/info/12/lacp/dump` command to check whether the ports are trunked.

NOTE: If you configure LACP on ports with 802.1x network access control, make sure the ports on both sides of the connection are properly configured for both LACP and 802.1x.

Configuring LACP

Use the following procedure to configure LACP for port 20 and port 21 to participate in link aggregation.

1. Set the LACP mode on port 20.

```
>> # /cfg/l2/lacp/port 20          (Select port 20)
>> LACP port 20# mode active      (Set port 20 to LACP active mode)
```

2. Define the admin key on port 20. Only ports with the same admin key can form a LACP trunk group.

```
>> LACP port 20# adminkey 100      (Set port 20 adminkey to 100)
Current LACP port adminkey:      17
New pending LACP port adminkey: 100
```

3. Set the LACP mode on port 21.

```
>> # /cfg/l2/lacp/port 21          (Select port 21)
>> LACP port 21# mode active      (Set port 21 to LACP active mode)
```

4. Define the admin key on port 21.

```
>> LACP port 21# adminkey 100      (Set port 21 adminkey to 100)
Current LACP port adminkey:      18
New pending LACP port adminkey: 100
```

5. Apply and verify the configuration.

```
>> LACP port 21# apply             (Make your changes active)
>> LACP port 21# cur              (View current trunking configuration)
```

6. Save your new configuration changes.

```
>> LACP port 21# save             (Save for restore after reboot)
```

Port-based Network Access and traffic control

Port-based Network Access control

Port-based Network Access control provides a means of authenticating and authorizing devices attached to a LAN port that has point-to-point connection characteristics. It prevents access to ports that fail authentication and authorization. This feature provides security to all ports of the HP 10GbE switch (except the management port 17).

The following topics are discussed in this section:

- Extensible Authentication Protocol over LAN
- 802.1x Authentication Process
- 802.1x Port States
- Supported RADIUS Attributes
- Configuration Guidelines

Extensible authentication protocol over LAN

HP 10GbE switch software can provide user-level security for its ports using the IEEE 802.1x protocol, which is a more secure alternative to other methods of port-based network access control. Any device attached to an 802.1x-enabled port that fails authentication is prevented access to the network and denied services offered through that port.

The 802.1x standard describes port-based network access control using Extensible Authentication Protocol over LAN (EAPoL). EAPoL provides a means of authenticating and authorizing devices attached to a LAN port that has point-to-point connection characteristics and of preventing access to that port in cases of authentication and authorization failures.

EAPoL is a client-server protocol that has the following components:

- **Supplicant or Client**—The Supplicant is a device that requests network access and provides the required credentials (user name and password) to the Authenticator and the Authentication Server.
- **Authenticator**—The Authenticator enforces authentication and controls access to the network. The Authenticator grants network access based on the information provided by the Supplicant and the response from the Authentication Server. The Authenticator acts as an intermediary between the Supplicant and the Authentication Server: requesting identity information from the client, forwarding that information (encapsulated in RADIUS packets) to the Authentication Server for validation, relaying the server's responses to the client, and authorizing network access based on the results of the authentication exchange. The HP 10GbE switch acts as an Authenticator.
- **Authentication Server**—The Authentication Server validates the credentials provided by the Supplicant to determine if the Authenticator should grant access to the network. The Authentication Server may be co-located with the Authenticator. The switch relies on external RADIUS servers for authentication.

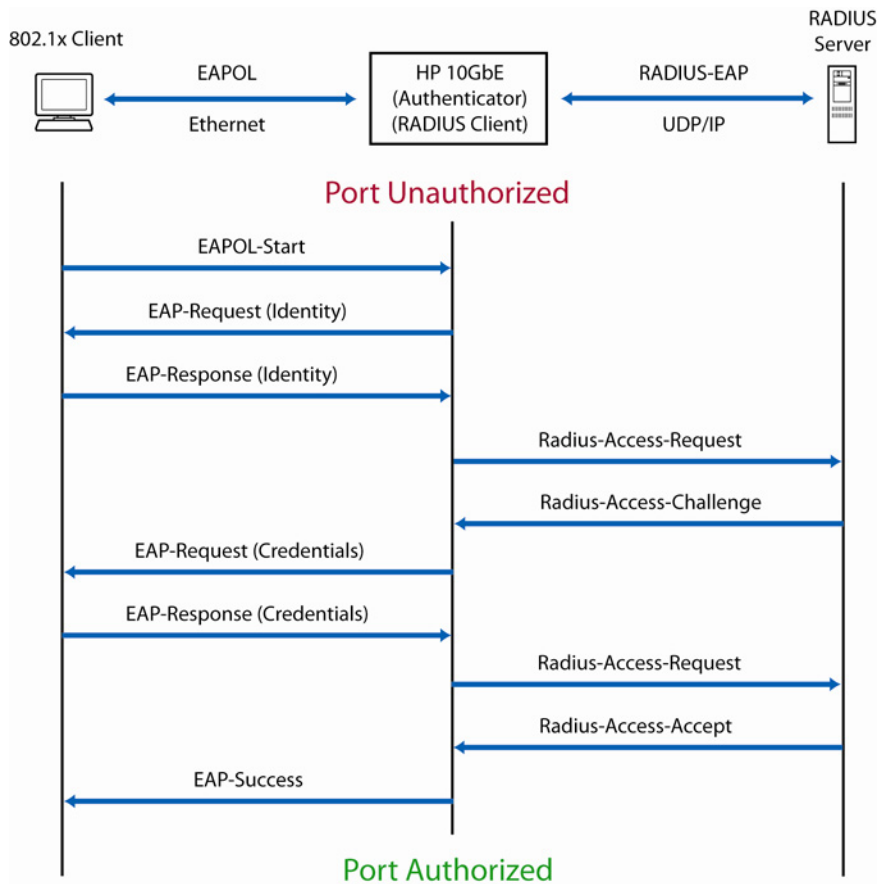
Upon a successful authentication of the client by the server, the 802.1x-controlled port transitions from unauthorized to authorized state, and the client is allowed full access to services through the port. When the client sends an EAP-Logoff message to the authenticator, the port will transition from authorized to unauthorized state.

802.1x authentication process

The clients and authenticators communicate using Extensible Authentication Protocol (EAP), which was originally designed to run over PPP, and for which the IEEE 802.1x Standard has defined an encapsulation method over Ethernet frames, called EAP over LAN (EAPoL).

The following figure shows a typical message exchange initiated by the client.

Figure 2 Using EAPoL to authenticate a port



EAPoL Message Exchange

During authentication, EAPoL messages are exchanged between the client and the switch authenticator, while RADIUS-EAP messages are exchanged between the switch authenticator and the Radius authentication server.

Authentication is initiated by one of the following methods:

Switch authenticator sends an EAP-Request/Identity packet to the client.

Client sends an EAPoL-Start frame to the switch authenticator, which responds with an EAP-Request/Identity frame.

The client confirms its identity by sending an EAP-Response/Identity frame to the switch authenticator, which forwards the frame encapsulated in a RADIUS packet to the server.

The Radius server chooses an EAP-supported authentication algorithm to verify the client's identity, and sends an EAP-Request packet to the client via the switch authenticator. The client then replies to the Radius server with an EAP-Response containing its credentials.

Upon a successful authentication of the client by the server, the 802.1x-controlled port transitions from unauthorized to authorized state, and the client is allowed full access to services through the controlled port. When the client later sends an EAPOL-Logoff message to the switch authenticator, the port transitions from authorized to unauthorized state.

If a client that does not support 802.1x connects to an 802.1x-controlled port, the switch authenticator requests the client's identity when it detects a change in the operational state of the port. The client does not respond to the request, and the port remains in the unauthorized state.

NOTE: When an 802.1x-enabled client connects to a port that is not 802.1x-controlled, the client initiates the authentication process by sending an EAPOL-Start frame. When no response is received, the client retransmits the request for a fixed number of times. If no response is received, the client assumes the port is in authorized state, and begins sending frames, even if the port is unauthorized.

802.1x port states

The state of the port determines whether the client is granted access to the network, as follows:

- **Unauthorized**—While in this state, the port discards all ingress and egress traffic except EAP packets.
- **Authorized**—When the client is authenticated successfully, the port transitions to the authorized state allowing all traffic to and from the client to flow normally.
- **Force Unauthorized**—You can configure this state that denies all access to the port.
- **Force Authorized**—You can configure this state that allows full access to the port.

Use the 802.1x Global Configuration Menu (`/cfg/12/8021x/global`) to configure 802.1x authentication for all ports in the switch. Use the 802.1x Port Menu (`/cfg/12/8021x/port x`) to configure a single port.

Supported RADIUS attributes

The HP 10GbE switch 802.1x Authenticator relies on external RADIUS servers for authentication with EAP. The following table lists the RADIUS attributes that are supported as part of RADIUS-EAP authentication based on the guidelines specified in Annex D of the 802.1x standard and RFC 3580.

Table 9 EAP support for RADIUS attributes

#	Attribute	Attribute Value	A-R	A-A	A-C	A-R
1	User-Name	The value of the Type-Data field from the supplicant's EAP-Response/Identity message. If the Identity is unknown (i.e. Type-Data field is zero bytes in length), this attribute will have the same value as the Calling-Station-Id.	1	0-1	0	0
4	NAS-IP-Address	IP address of the authenticator used for RADIUS communication.	1	0	0	0
5	NAS-Port	Port number of the authenticator port to which the supplicant is attached.	1	0	0	0
24	State	Server-specific value. This is sent unmodified back to the server in an Access-Request that is in response to an Access-Challenge.	0-1	0-1	0-1	0
30	Called-Station-ID	The MAC address of the authenticator encoded as an ASCII string in canonical format, e.g. 000D5622E3 9F.	1	0	0	0
31	Calling-Station-ID	The MAC address of the supplicant encoded as an ASCII string in canonical format, e.g. 00034B436206.	1	0	0	0
79	EAP-Message	Encapsulated EAP packets from the supplicant to the authentication server (Radius) and vice-versa. The authenticator relays the decoded packet to both devices.	1+	1+	1+	1+
80	Message-Authenticator	Always present whenever an EAP-Message attribute is also included. Used to integrity-protect a packet.	1	1	1	1
87	NAS-Port-ID	Name assigned to the authenticator port, e.g. Server1_Port3	1	0	0	0

Legend:

RADIUS Packet Types: A-R (Access-Request), A-A (Access-Accept), A-C (Access-Challenge), A-R (Access-Reject)

RADIUS Attribute Support:

0—This attribute MUST NOT be present in a packet.

0+—Zero or more instances of this attribute MAY be present in a packet.

0-1—Zero or one instance of this attribute MAY be present in a packet.

1—Exactly one instance of this attribute MUST be present in a packet.

1+—One or more of these attributes MUST be present.

EAPoL configuration guidelines

When configuring EAPoL, consider the following guidelines:

- The 802.1x port-based authentication is currently supported only in point-to-point configurations, that is, with a single supplicant connected to an 802.1x-enabled switch port.
- When 802.1x is enabled, a port has to be in the authorized state before any other Layer 2 feature can be operationally enabled. For example, the STG state of a port is operationally disabled while the port is in the unauthorized state.
- The 802.1x supplicant capability is not supported. Therefore, none of its ports can connect successfully to an 802.1x-enabled port of another device, such as another switch, which acts as an authenticator, unless access control on the remote port is disabled or is configured in forced-authorized mode. For example, if a HP 10GbE switch is connected to another HP 10GbE switch, and if 802.1x is enabled on both switches, the two connected ports must be configured in force-authorized mode.
- The 802.1x standard has optional provisions for supporting dynamic virtual LAN assignment via RADIUS tunneling attributes, for example, Tunnel-Type (=VLAN), Tunnel-Medium-Type (=802), and Tunnel-Private-Group-ID (=VLAN id). These attributes are not supported and might affect 802.1x operations. Other unsupported attributes include Service-Type, Session-Timeout, and Termination-Action.

RADIUS accounting service for 802.1x-authenticated devices or users is not supported.

Configuration changes performed using SNMP and the standard 802.1x MIB take effect immediately.

Port-based traffic control

Port-based traffic control prevents HP 10GbE switch ports from being disrupted by LAN storms. A LAN storm occurs when data packets flood the LAN, which can cause the network to become congested and slow down. Errors in the protocol-stack implementation or in the network configuration can cause a LAN storm.

You can enable port-based traffic control separately for each of the following traffic types:

- Broadcast—packets with destination MAC address ff:ff:ff:ff:ff:ff
- Multicast—packets that have MAC addresses with the least significant bit of their first octet set to one
- Destination Lookup Failed (DLF)—packets with unknown destination MAC address, that are treated like broadcast packets

With Port-based Traffic Control enabled, the port monitors incoming traffic of each type noted above. If the traffic exceeds a configured threshold, the port blocks traffic that exceeds the threshold until the traffic flow falls back within the threshold.

The HP 10GbE switch supports separate traffic-control thresholds for broadcast, multicast, and DLF traffic. The traffic threshold is measured in number of frames per second.

NOTE: All ports that belong to a trunk must have the same traffic-control settings.

Configuring port-based traffic control

To configure a port for traffic control, perform the following steps:

1. Configure the traffic-control threshold and enable traffic control.

```
Main# /cfg/port 2
>> Port 2# brate 150000          (Set broadcast threshold)
>> Port 2# mrate 150000        (Set multicast threshold)
>> Port 2# drate 150000        (Set DLF threshold)
```

2. To disable a traffic-control threshold, use the following command:

```
>> Port 2# mrate dis          (Disable multicast threshold)
```

3. Apply and save the configuration.

```
>> Port 2# apply              (Apply the port configurations)
>> Port 2# save                (Save the port configurations)
```

VLANs

Introduction

This chapter describes network design and topology considerations for using Virtual Local Area Networks (VLANs). VLANs are commonly used to split up groups of network users into manageable broadcast domains, to create logical segmentation of workgroups, and to enforce security policies among logical segments.

The following topics are discussed in this chapter:

- VLANs and Port VLAN ID Numbers
- VLAN Tagging
- VLANs and IP Interfaces
- VLAN Topologies and Design Considerations

NOTE: Basic VLANs can be configured during initial switch configuration.

More comprehensive VLAN configuration can be done from the command line interface. See the *HP 10Gb Ethernet BL-c Switch Command Reference Guide*.

Overview

Setting up VLANs is a way to segment networks to increase network flexibility without changing the physical network topology. With network segmentation, each switch port connects to a segment that is a single broadcast domain. When a switch port is configured to be a member of a VLAN, it is added to a group of ports (workgroup) that belongs to one broadcast domain.

Ports are grouped into broadcast domains by assigning them to the same VLAN. Multicast, broadcast, and unknown unicast frames are flooded only to ports in the same VLAN.

VLANs and port VLAN ID numbers

VLAN numbers

The HP 10GbE switch supports up to 1,000 VLANs per switch. Even though the maximum number of VLANs supported at any given time is 1,000, each can be identified with any number between 1 and 4095. VLAN 1 is the default VLAN, and all ports are assigned to it. VLAN 4095 is reserved for switch management, and it cannot be configured.

Viewing VLANs

The VLAN information menu (`/info/12/vlan`) displays all configured VLANs and all member ports that have an active link state, for example:

```
>> Layer 2# vlan
VLAN          Name                Status Ports
-----
1             Default VLAN        ena    1 4-16 18-21
2             VLAN 2              ena    2 3
4095          VLAN 4095           ena    17
```

PVID numbers

Each port in the switch has a configurable default VLAN number, known as its PVID. This places all ports on the same VLAN initially, although each port PVID is configurable to any VLAN number between 1 and 4094.

The default configuration settings for switches have all ports set as untagged members of VLAN 1 with all ports configured as PVID = 1. In the default configuration example shown in the following figure, all incoming packets are assigned to VLAN 1 by the default port VLAN identifier (PVID = 1).

Viewing and configuring PVIDs

You can view PVIDs from the following CLI commands:

Port information

```
>> /info/port
Port Tag RMON PVID      NAME                VLAN (s)
-----
1    n  d    1 Downlink1          1
2    n  e    1 Downlink2          1
3    n  d    1 Downlink3          1
4    n  d    1 Downlink4          1
5    n  d    1 Downlink5          1
6    n  d    1 Downlink6          1
7    n  d    1 Downlink7          1
:
:
```

Port configuration

```
>> /cfg/port 21/pvid 21
Current port VLAN ID:    1
New pending port VLAN ID: 21

>> Port 21#
```

Each port on the switch can belong to one or more VLANs, and each VLAN can have any number of switch ports in its membership. Any port that belongs to multiple VLANs, however, must have VLAN tagging enabled. See the “VLAN tagging” section in this chapter.

Any untagged frames (those with no VLAN specified) are classified with the PVID of the sending port.

VLAN tagging

The switch supports IEEE 802.1Q VLAN tagging, providing standards-based VLAN support for Ethernet systems.

Tagging places the VLAN identifier in the frame header, allowing each port to belong to multiple VLANs. When you configure multiple VLANs on a port, you must also enable tagging on that port.

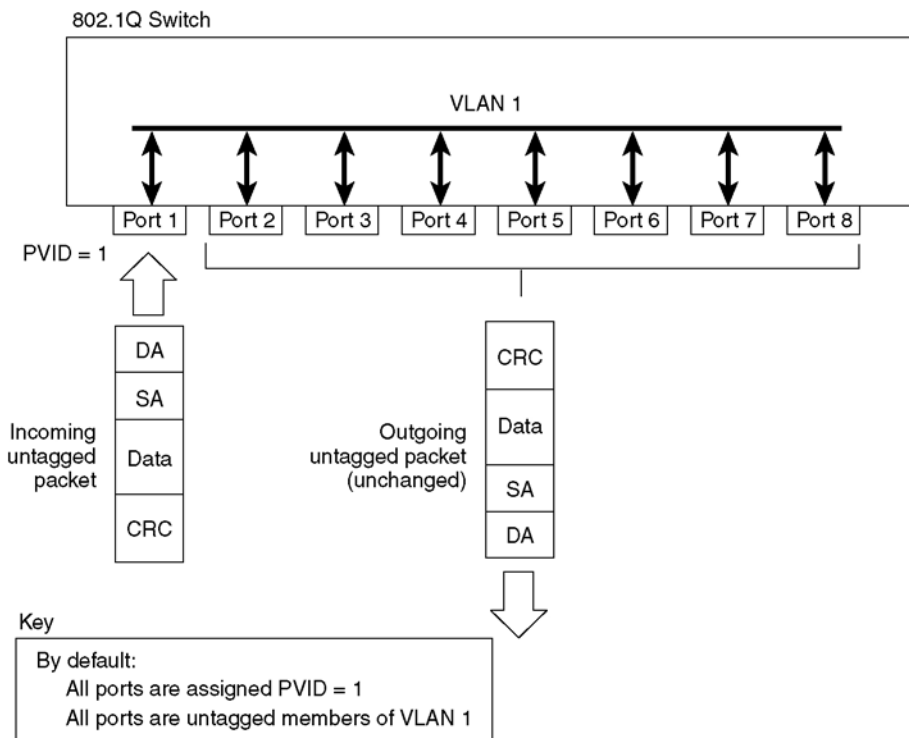
Since tagging fundamentally changes the format of frames transmitted on a tagged port, you must carefully plan network designs to prevent tagged frames from being transmitted to devices that do not support 802.1Q VLAN tags, or devices where tagging is not enabled.

Important terms used with the 802.1Q tagging feature are:

- VLAN identifier (VID)—the 12-bit portion of the VLAN tag in the frame header that identifies an explicit VLAN.
- Port VLAN identifier (PVID)—a classification mechanism that associates a port with a specific VLAN. For example, a port with a PVID of 3 (PVID = 3) assigns all untagged frames received on this port to VLAN 3.
- Tagged frame—a frame that carries VLAN tagging information in the header. The VLAN tagging information is a 32-bit field (VLAN tag) in the frame header that identifies the frame as belonging to a specific VLAN. Untagged frames are marked (tagged) with this classification as they leave the switch through a port that is configured as a tagged port.
- Untagged frame—a frame that does not carry any VLAN tagging information in the frame header.
- Untagged member—a port that has been configured as an untagged member of a specific VLAN. When an untagged frame exits the switch through an untagged member port, the frame header remains unchanged. When a tagged frame exits the switch through an untagged member port, the tag is stripped and the tagged frame is changed to an untagged frame.
- Tagged member—a port that has been configured as a tagged member of a specific VLAN. When an untagged frame exits the switch through a tagged member port, the frame header is modified to include the 32-bit tag associated with the PVID. When a tagged frame exits the switch through a tagged member port, the frame header remains unchanged (original VID remains).

NOTE: If an 802.1Q tagged frame is sent to a port that has VLAN-tagging disabled, then the frames are forwarded based on their port-VLAN ID (PVID).

Figure 3 Default VLAN settings



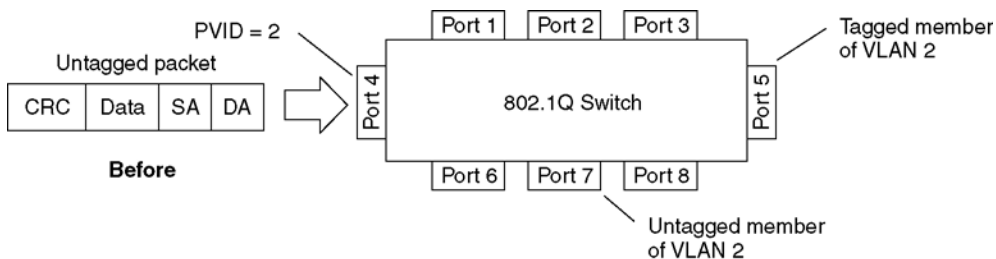
BS45010A

NOTE: The port numbers specified in these illustrations may not directly correspond to the physical port configuration of your switch model.

When you configure VLANs, you configure the switch ports as tagged or untagged members of specific VLANs. See the following figures.

In the following figure, the untagged incoming packet is assigned directly to VLAN 2 (PVID = 2). Port 5 is configured as a tagged member of VLAN 2, and port 7 is configured as an untagged member of VLAN 2.

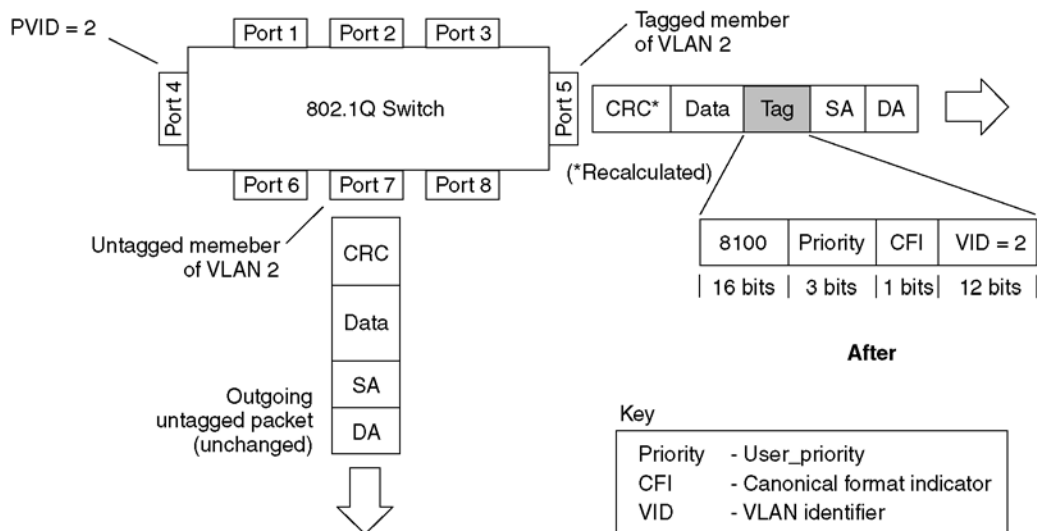
Figure 4 Port-based VLAN assignment



BS45011A

As shown in the following figure, the untagged packet is marked (tagged) as it leaves the switch through port 5, which is configured as a tagged member of VLAN 2. The untagged packet remains unchanged as it leaves the switch through port 7, which is configured as an untagged member of VLAN 2.


Figure 5 802.1Q tagging (after port-based VLAN assignment)



BS45012A


In the following figure, the tagged incoming packet is assigned directly to VLAN 2 because of the tag assignment in the packet. Port 5 is configured as a tagged member of VLAN 2, and port 7 is configured as an untagged member of VLAN 2.

Figure 6 802.1Q tag assignment



BS45013A

As shown in the following figure, the tagged packet remains unchanged as it leaves the switch through port 5, which is configured as a tagged member of VLAN 2. However, the tagged packet is stripped (untagged) as it leaves the switch through port 7, which is configured as an untagged member of VLAN 2.

Figure 7 802.1Q tagging (after 802.1Q tag assignment)

NOTE: Using the `/boot/conf factory` command resets all ports to VLAN 1 (except management port 17) and all other settings to the factory defaults at the next reboot.

VLANs and IP interfaces

Carefully consider how you create VLANs within the switch, so that communication with the switch remains possible. In order to access the switch for remote configuration, trap messages, and other management functions, be sure that at least one IP interface on the switch has a VLAN defined.

You can also inadvertently cut off access to management functions if you exclude the ports from the VLAN membership. For example, if all IP interfaces are left on VLAN 1 (the default), and all ports are configured for VLAN 2, and then switch management features are effectively cut off.

To remedy this, keep all ports used for remote switch management on the default VLAN and assign an IP interface to the default VLAN.

For more information on configuring IP interfaces, see the “Configuring an IP interface” section in the “Accessing the switch” chapter.

VLAN topologies and design considerations

By default, all switch ports are configured to the default VLAN 1. This configuration groups all ports into the same broadcast domain. The VLAN has an 802.1Q VLAN ID of 1. VLAN tagging is turned off, because, by default, all ports are members of a single VLAN only.

If configuring Spanning Tree Protocol (`/cfg/12/stp`), note that each of spanning tree groups 2-128 may contain only one VLAN. If configuring Multiple Spanning Tree Protocol (`/cfg/12/mrst`), each of the spanning tree groups (1-32 for MSTP) may contain multiple VLANs.

VLAN configuration rules


VLANs operate according to specific configuration rules which must be considered when creating VLANs. For example:

- HP recommends that all ports involved in trunking and Port Mirroring have the same VLAN configuration. If a port is on a trunk with a mirroring port, the VLAN configuration cannot be changed. For more information on port trunking, see the “Port trunking example” section in the “Ports and trunking” chapter.
- All ports involved in Port Mirroring must have memberships in the same STP group, but they can have memberships in different VLANs.
- When you delete a VLAN, untagged ports are moved to the default VLAN (VLAN 1). Tagged ports that belong only to the deleted VLAN are moved to the default VLAN 1. Tagged ports that belong to multiple VLANs are removed from the deleted VLAN only.

Multiple VLANs with tagging

The following figure shows only those switch-port-to-server links that must be configured for the example. While not shown, all other server links remain set at their default settings.

Figure 8 Multiple VLANs with VLAN tagging



The features of this VLAN are described in the following table:

Table 10 Multiple VLANs with tagging

Component	Description
Switch 1	Switch 1 is configured for VLANs 1, 2, and 3. Port 1 is tagged to accept traffic from VLANs 1 and 2. Port 18 is tagged to accept traffic from VLANs 1, 2, and 3. Port 20 is an untagged member of VLAN 2. Port 21 is tagged, and accepts traffic from VLANs 1 and 3.
Switch 2	Switch 2 is configured for VLANs 1, 3, and 4. Port 2 is tagged to accept traffic from VLANs 3 and 4. Port 4 is configured only for VLAN 3, so VLAN tagging is off. Port 18 is tagged to accept traffic from VLANs 1 and 3. Port 20 is an untagged member of VLAN 4. Port 21 is tagged, and accepts traffic from VLANs 1 and 3.

Table 10 Multiple VLANs with tagging

Component	Description
Blade Server #1	This high-use blade server needs to be accessed from all VLANs and IP subnets. The server has a VLAN-tagging adapter installed with VLAN tagging turned on. One adapter is attached to one of the switch's 10 Gbps ports, that is configured for VLANs 1 and 2. One adapter is configured for VLANs 3 and 4. Because of the VLAN tagging capabilities of both the adapter and the switch, the server is able to communicate on all four VLANs in this network while maintaining broadcast separation among all four VLANs and subnets.
Blade Server #2	This blade server belongs to VLAN 3. The port that the VLAN is attached to is configured only for VLAN 3, so VLAN tagging is off.
PC #1	This PC is a member of VLAN 2 and 3. Via VLAN 2, it can communicate with Server 1, PC 3, and PC 5. Using VLAN 3, it can communicate with Server 1, Server 2, and PC 4.
PC #2	This PC is a member of VLAN 4, and can only communicate with Server 1.
PC #3	This PC is a member of VLAN 1 and VLAN 2. Using VLAN 1, it can communicate with Server 1 and PC 5. Using VLAN 2, it can communicate with Server 1, PC 1, and PC 5.
PC #4	This PC is a member of VLAN 3, and it can communicate with Server 1, Server 2, and PC 1.
PC #5	This PC is a member of both VLAN 1 and VLAN 2. Using VLAN 1, it can communicate with Server 1 and PC 3. Using VLAN 2, it can communicate with Server 1, PC 1, and PC 3. The Layer 2 switch port to which it is connected is configured for both VLAN 1 and VLAN 2 and has tagging enabled.

NOTE: All PCs connected to a tagged port must have an Ethernet adapter with VLAN-tagging capability installed.

Configuring the example network

These examples describe how to configure ports and VLANs on Switch 1 and Switch 2.

Configuring ports and VLANs on Switch 1 (CLI example)

To configure ports and VLANs on Switch 1, do the following:

1. On Switch 1, enable VLAN tagging on the necessary ports.

```

Main# /cfg/port 1
>> Port 1# tag e                               (Select port 1: connection to server 1)

Current VLAN tag support: disabled
New VLAN tag support:    enabled              (Enable tagging)
Port 1 changed to tagged.

Main# /cfg/port 18
>> Port 18# tag e                               (Select uplink port 18)
                                         (Enable tagging)
Current VLAN tag support: disabled
New VLAN tag support:    enabled
Port 18 changed to tagged.
>> Port 18# apply                               (Apply the port configurations)

```

2. Configure the VLANs and their member ports. Since all ports are by default configured for VLAN 1, configure only those ports that belong to VLAN 2.

```
>> /cfg/l2/vlan 2

>> VLAN 2# add 1                                (Add port 1 to VLAN 2)
Current ports for VLAN 2: empty
Pending new ports for VLAN 2: 1

>> VLAN 2# add 18                                (Add port 18 to VLAN 2)
Current ports for VLAN 2: 1
Pending new ports for VLAN 2: 18

>> VLAN 2# add 20                                (Add port 20 to VLAN 2)
Port 20 is an UNTAGGED port and its current PVID is 1.
Confirm changing PVID from 1 to 2 [y/n]: y
Current ports for VLAN 2: 1, 18
Pending new ports for VLAN 2: 20

>> VLAN 3# add 18                                (Add port 18 to VLAN 3)
Current ports for VLAN 3:
Pending new ports for VLAN 3: 18

>> apply                                          (Apply the port configurations)
>> save                                          (Save the port configurations)
```

Configuring ports and VLANs on Switch 2 (CLI example)

To configure ports and VLANs on Switch 2, do the following:

1. On Switch 2, enable VLAN tagging on the necessary ports. Port 4 (connection to server 2) remains untagged, so it is not configured below.

```

Main# /cfg/port 2                               (Select port 2: connection to server 1)
>> Port 2# tag e
Current VLAN tag support: disabled
New VLAN tag support:     enabled
Port 2 changed to tagged.

Main# /cfg/port 18                              (Select uplink port 18)
>> Port 18# tag e                               (Enable tagging)
Current VLAN tag support: disabled
New VLAN tag support:     enabled
Port 18 changed to tagged.

>> Port 18# apply                               (Apply the port configurations)

```

2. Configure the VLANs and their member ports. Since all ports are by default configured for VLAN 1, configure only those ports that belong to other VLANs.

```

>> /cfg/l2/vlan 3
>> VLAN 3# add 2
Current ports for VLAN 3: empty
Pending new ports for VLAN 3: 2

>> VLAN 3# add 4
Port 4 is an UNTAGGED port and its current PVID is 1.
Confirm changing PVID from 1 to 2 [y/n]: y
Current ports for VLAN 3: 2
Pending new ports for VLAN 3:

>> VLAN 3# add 18
Current ports for VLAN 3: 2
Pending new ports for VLAN 3: 18

>> /cfg/l2/vlan 4
>> VLAN 4# add 20
Port 20 is an UNTAGGED port and its current PVID is 1.
Confirm changing PVID from 1 to 2 [y/n]: y
Current ports for VLAN 4: empty
Pending new ports for VLAN 4: 20

>> apply                                       (Apply the port configurations)
>> save                                       (Save the port configurations)

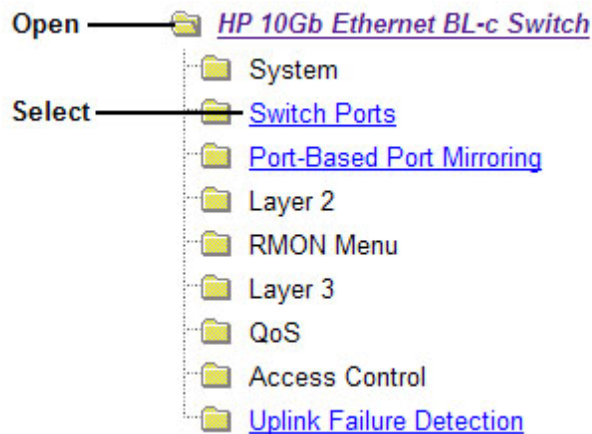
```

The external layer 2 switches should also be configured for VLANs and tagging.

Configuring ports and VLANs on Switch 1 (BBI example)

To configure ports and VLANs on Switch 1, do the following:

1. On the switch 1, enable VLAN tagging on the necessary ports.
 - a. Click the **Configure** context button on the Toolbar.
 - b. Open the Switch folder, and select **Switch Ports** (click the underlined text, not the folder).



- c. Click a port number to select it.


Switch Ports Configuration								
Switch Port	State	VLAN Tagging	Default PVID	PVID tagging	Multicast Threshold	Broadcast Threshold	Destination Lookup Fail Threshold	802.1p Priority
Select — <u>1</u>	disabled	disabled	1	enabled	disabled	disabled	disabled	0
<u>2</u>	disabled	disabled	1	enabled	disabled	disabled	disabled	0
<u>3</u>	disabled	disabled	1	enabled	disabled	disabled	disabled	0
<u>4</u>	disabled	disabled	1	enabled	disabled	disabled	disabled	0
<u>5</u>	disabled	disabled	1	enabled	disabled	disabled	disabled	0
<u>6</u>	disabled	disabled	1	enabled	disabled	disabled	disabled	0
<u>7</u>	disabled	disabled	1	enabled	disabled	disabled	disabled	0
<u>8</u>	disabled	disabled	1	enabled	disabled	disabled	disabled	0
<u>9</u>	disabled	disabled	1	enabled	disabled	disabled	disabled	0
<u>10</u>	disabled	disabled	1	enabled	disabled	disabled	disabled	0

- d. Enable the port and enable VLAN tagging.

Switch Port 1 Configuration

Switch Port State	Enabled ▾
RMON Instrumentation	Disabled ▾
VLAN Tagging	Enabled ▾
PVID Tagging	Disabled ▾
Port STP	Off ▾
Default Port VLAN ID (1 - 4095)	1

- e. Click **Submit**.
2. Configure the VLANs and their member ports.
- a. Open the Virtual LANs folder, and select **Add VLAN**.



- b. Enter the VLAN name, VLAN ID number, and enable the VLAN. To add ports, select each port in the Ports Available list and click **Add**. Since all ports are configured for VLAN 1 by default, configure only those ports that belong to VLAN 2.

VLAN "New" Configuration

VLAN Name	VLAN 2
VLAN ID (1 - 4095) From	2
VLAN State	enabled ▼
Spanning Tree Group	2

Ports Available

PortID ▲
 2
3
4
5
6
7
8
9
10
▼

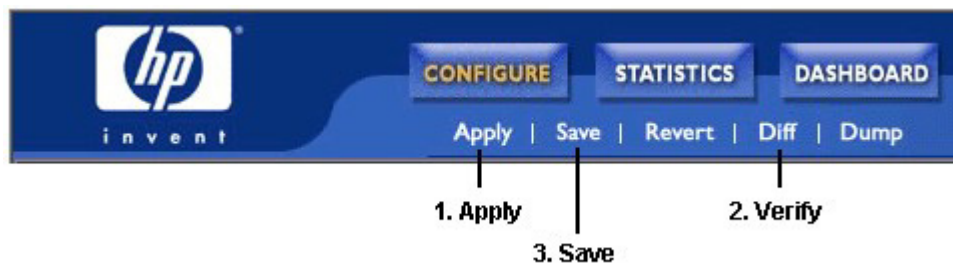
Ports in Vlan

PortID ▲
 1
18
20
▼

- c. Click **Submit**.

The external Layer 2 switches should also be configured for VLANs and tagging.

3. Apply, verify, and save the configuration.



FDB static entries

Static entries in the Forwarding Database (FDB) allow the switch to forward packets without flooding ports to perform a lookup. A FDB static entry is a MAC address associated with a specific port and VLAN. The switch supports 128 static entries. Static entries are manually configured, using the following command:

```
/cfg/l2/fdb/static
```

FDB static entries are permanent, so the FDB Aging value does not apply to them. Static entries are manually added to the FDB, and manually deleted from the FDB.

Incoming frames that contain the static entry as the source MAC can use only ports configured for the static entry.

Trunking support for FDB static entries

A FDB static entry can be added to a port that is a member of a trunk group, as follows:

- Static (manually configured) trunk group
- Dynamic (LACP) trunk group

The trunk group supports the FDB static entry. If the port with the static entry fails, other ports in the trunk handle the traffic. If the port is removed from the trunk, the static entry is removed from the trunk, but remains configured on the port.

The FDB information commands (`/info/l2/fdb`) display trunk support for static FDB entries, if applicable:

```
>> Forwarding Database# dump
      MAC address      VLAN  Port  Trnk  State
      -----
00:00:2e:9b:db:f8     1      1      1    TRK
00:00:5e:00:01:f4     1     21      1    FWD
00:01:81:2e:b5:60     1     21      1    FWD
00:02:a5:e9:76:30     1      1      1    TRK
00:03:4b:e2:15:f1     1     21      1    FWD
```

Configuring a static FDB entry

Perform the following actions to configure a static FDB entry:

```
Main# /cfg/l2/fdb/static (Select static FDB menu)
>> Static FDB# add 00:60:af:00:02:30
Enter VLAN number: 2
Enter port (1-21): 2
>> Static FDB# apply (Apply the configuration)
>> Static FDB# save (Save the configuration)
```

Spanning Tree Protocol

Introduction

When multiple paths exist on a network, Spanning Tree Protocol (STP) configures the network so that a switch uses only the most efficient path. The following topics are discussed in this chapter:

- Overview
- Bridge Protocol Data Units (BPDUs)
- Spanning Tree Group (STG) configuration guidelines
- Multiple Spanning Trees

Overview

Spanning Tree Protocol (STP) detects and eliminates logical loops in a bridged or switched network. STP forces redundant data paths into a standby (blocked) state. When multiple paths exist, STP configures the network so that a switch uses only the most efficient path. If that path fails, STP automatically sets up another active path on the network to sustain network operations.

The switch supports IEEE 802.1d Spanning Tree Protocol for STG 1, and Per VLAN Spanning Tree Protocol (PVST+) for STGs 2-128, by default.

NOTE: The switch also supports IEEE 802.1w Rapid Spanning Tree Protocol and IEEE 802.1s Multiple Spanning Tree Protocol. For more information, see the “RSTP and MSTP” chapter in this guide.

Bridge Protocol Data Units

To create a spanning tree, the application switch generates a configuration Bridge Protocol Data Unit (BPDU), which it then forwards out of its ports. All switches in the Layer 2 network participating in the spanning tree gather information about other switches in the network through an exchange of BPDUs.

A BPDU is a 64-byte packet that is sent out at a configurable interval, which is typically set for two seconds. The BPDU is used to establish a path, much like a `hello` packet in IP routing. BPDUs contain information about the transmitting bridge and its ports, including bridge and MAC addresses, bridge priority, port priority, and port path cost. If the ports are tagged, each port sends out a special BPDU containing the tagged information.

The generic action of a switch on receiving a BPDU is to compare the received BPDU to its own BPDU that it will transmit. If the received BPDU has a priority value closer to zero than its own BPDU, it will replace its BPDU with the received BPDU. Then, the application switch adds its own bridge ID number and increments the path cost of the BPDU. The application switch uses this information to block any redundant paths.

Determining the path for forwarding BPDUs

When determining which port to use for forwarding and which port to block, the switch uses information in the BPDU, including each bridge priority ID. A technique based on the lowest root cost is then computed to determine the most efficient path for forwarding.

Bridge priority

The bridge priority parameter controls which bridge on the network is the STP root bridge. To make one switch the root bridge, configure the bridge priority lower than all other switches and bridges on your network. The lower the value, the higher the bridge priority. The bridge priority is configured using the following command:

```
/cfg/l2/stp x/brg/prior
```

Port priority

The port priority helps determine which bridge port becomes the designated port. In a network topology that has multiple bridge ports connected to a single segment, the port with the lowest port priority becomes the designated port for the segment. The port priority is configured using the following command:

```
/cfg/l2/stp y/port x/prior
```

Port path cost

The port path cost assigns lower values to high-bandwidth ports, such as Gigabit Ethernet, to encourage their use. The objective is to use the fastest links so that the route with the lowest cost is chosen. A value of zero indicates that port cost is computed dynamically based on link speed. This works when forcing link speed, so it does not just apply to auto negotiated link speed.

By default, all switch ports have the path cost set to 2. To use dynamic path cost, based on link speed, set the path cost to 0 (zero). For example, if the path cost is set to zero:

- A 10 Gbps link receives a path cost of 2
- A 100 Mbps link receives a path cost of 19

Configure the port path cost using the following command:

```
/cfg/l2/stp y/port x/cost
```

Spanning Tree Group configuration guidelines

This section provides important information on configuring Spanning Tree Groups (STGs).

Default Spanning Tree configuration

In the default configuration, a single STG with the ID of 1 includes all ports on the switch. It is called the default STG. All other STGs (except the default STG) are empty, and VLANs must be added by the user.

You cannot assign ports directly to an STG. Add the ports to a VLAN, and add the VLAN to the STG. STGs 1-127 are enabled by default and assigned an ID number from 1 to 127. STG 128 is disabled by default, and contains the management VLAN 4095.

An STG cannot be deleted, only disabled. If you disable the STG while it still contains VLAN members, Spanning Tree will be off on all ports belonging to that VLAN.

Adding a VLAN to a Spanning Tree Group

If no VLANs exist beyond the default VLAN 1, see the “Creating a VLAN” section in this chapter for information on adding ports to VLANs.

Add the VLAN to the STG using the command `/cfg/12/stp <stg number>/add <vlan number>`.

Creating a VLAN

When you create a VLAN, then that VLAN automatically belongs to STG 1, the default STG. If you want the VLAN in another STG, you must move the VLAN by assigning it to another STG.

To move a newly created VLAN to an existing STG:

1. Create the VLAN.
2. Add the VLAN to an existing STG.

When creating a VLAN also consider the following:

- A VLAN cannot belong to more than one STG.
- VLANs that span multiple switches must be mapped within the same Spanning Tree Group (have the same STG ID) across all the switches.

Rules for VLAN tagged ports

Rules for VLAN tagged ports are listed below:

- If a port is tagged, it can belong to multiple STGs.
- When a tagged port belongs to more than one STG, the egress BPDUs are tagged to distinguish the BPDUs of one STG from those of another STG.
- An untagged port cannot span multiple STGs.

Adding and removing ports from STGs

Information on adding and removing ports from STGs is as follows:

- By default, all ports belong to VLAN 1 and STG 1.
- Each port is always a member of at least one VLAN. Each VLAN is always a member of at least one STG. Port membership within VLANs can be changed, and VLAN membership within STGs can be changed. To move a port from one STG to another, move the VLAN to which the port belongs, or move the port to a VLAN that belongs to the STG.
- When you remove a port from a VLAN, that port is also removed from the STG to which the VLAN belongs. However, if that port belongs to another VLAN in the same STG, the port remains in the STG.
- If you remove an untagged port from a non-default VLAN and STG, it is added to VLAN 1 and STG 1.

The relationship between ports, trunk groups, VLANs, and spanning trees is shown in the following table.

Table 11 Ports, trunk groups, and VLANs

Switch element	Belongs to
Port	Trunk group, or one or more VLANs
Trunk group	Only one VLAN
VLAN (non-default)	One Spanning Tree Group

Assigning cost to ports and trunk groups

When you configure a trunk group to participate in a Spanning Tree Group, all ports must have the same Spanning Tree configuration, as follows:

- port priority
- path cost
- link type
- Edge port status
- Port Fast Forward status

Assign lower path costs on each member of a trunk group, to ensure the trunk group remains in the Forwarding state.

Multiple Spanning Trees

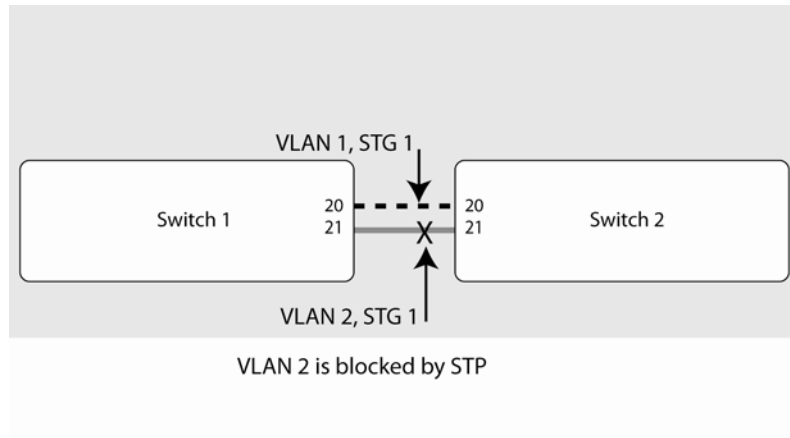
Each switch supports a maximum of 128 Spanning Tree Groups (STGs). Multiple STGs provide multiple data paths, which can be used for load-balancing and redundancy.

You enable independent links on two switches using multiple STGs by configuring each path with a different VLAN and then assigning each VLAN to a separate STG. Each STG is independent. Each STG sends its own Bridge Protocol Data Units (BPDUs), and each STG must be independently configured.

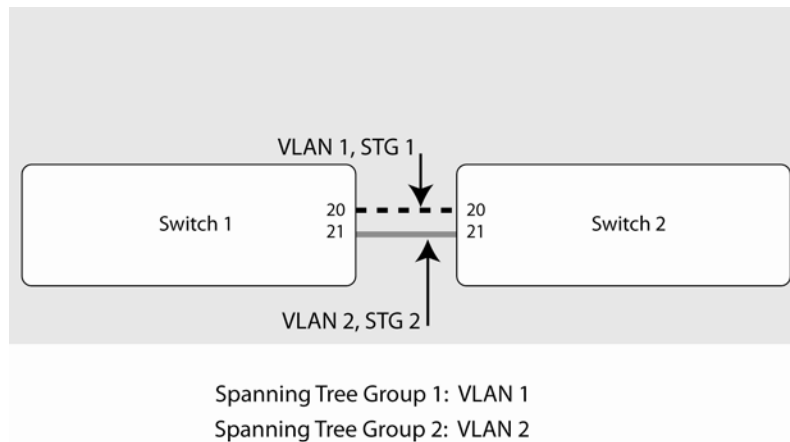
The STG, or bridge group, forms a loop-free topology that includes one or more virtual LANs (VLANs). The switch supports 128 STGs running simultaneously. The default STG 1 supports IEEE 802.1d Spanning Tree Protocol, and may contain more than one VLAN. All other STGs support Per VLAN Spanning Tree (PVST+), and may contain only one VLAN each. The switch can support multiple VLANs in STGs 2-128; however, you must enable IEEE 802.1s Multiple Spanning Tree Protocol mode. For more information, see the “RSTP and MSTP” chapter in this guide.

Why do we need Multiple Spanning Trees?

The following figure shows a simple example of why we need multiple Spanning Trees. This example assumes that port 20 and 21 are not part of a Trunk Group. Two VLANs (VLAN 1 and VLAN 2) exist between Switch 1 and Switch 2. If the same Spanning Tree Group is enabled on both switches, the switches see an apparent loop and block port 21 on Switch 2, which cuts off communication between the switches for VLAN 2.

Figure 9 Two VLANs on one instance of Spanning Tree Protocol

In the following figure, VLAN 1 and VLAN 2 belong to different Spanning Tree Groups. The two instances of spanning tree separate the topology without forming a loop, so that both VLANs can forward packets between the switches without losing connectivity.

Figure 10 Two VLANs on separate instances of Spanning Tree Protocol

VLAN participation in Spanning Tree Groups

The following table shows which switch ports participate in each Spanning Tree Group. By default, server ports (ports 1-16) do not participate in Spanning Tree, even though they are members of their respective VLANs.

Table 12 VLAN participation in Spanning Tree Groups

	VLAN 1	VLAN 2
Switch 1	Spanning Tree Group 1 Port 20	Spanning Tree Group 2 Port 21
Switch 2	Spanning Tree Group 1 Port 20	Spanning Tree Group 2 Port 21

Configuring Multiple Spanning Tree Groups

This section explains how to assign each VLAN to its own Spanning Tree Group on the switches 1 and 2.

By default, Spanning Tree Groups 2-127 are empty, and Spanning Tree Group 1 contains all configured VLANs until individual VLANs are explicitly assigned to other Spanning Tree Groups. Except for the default Spanning Tree Group 1, which may contain more than one VLAN, Spanning Tree Groups 2-128 may contain only one VLAN each.

NOTE: Each instance of Spanning Tree Group is enabled by default.

Configuring Switch 1 (CLI example)

1. Configure port and VLAN membership on Switch 1 as described in the “Configuring ports and VLANs on Switch 1 (CLI example)” section, in the “VLANs” chapter of this guide.
2. Add VLAN 2 to Spanning Tree Group 2.

```
>> /cfg/12/stp 2 (Select Spanning Tree Group 2)
>> Spanning Tree Group 2# add 2 (Add VLAN 2)
```

VLAN 2 is removed from Spanning Tree Group 1.

3. Apply and save.

```
>> apply (Apply the port configurations)
>> save (Save the port configurations)
```

Configuring Switch 2 (CLI example)

1. Configure port and VLAN membership as described in the “Configuring ports and VLANs on Switch 2 (CLI example)” section in the “VLANs” chapter of this guide.
2. Add VLAN 2 to Spanning Tree Group 2.

```
>> /cfg/12/stp 2 (Select Spanning Tree Group 2)
>> Spanning Tree Group 2# add 2 (Add VLAN 2)
```


VLAN 2 is removed from Spanning Tree Group 1.

3. Apply and save.

```
>> apply (Apply the port configurations)
>> save (Save the port configurations)
```

Configuring Switch 1 (BBI example)

1. Configure port and VLAN membership on Switch 1 as described in the “Configuring ports and VLANs on Switch 1 (BBI example)” section, in the “VLANs” chapter of this guide.
2. Add VLAN 2 to Spanning Tree Group 2.
 - a. Click the **Configure** context button on the Toolbar.
 - b. Select **Spanning Tree Groups** (click the underlined text, not the folder).



- c. Select a Spanning Tree Group number.

Switch Spanning Tree Groups Configuration

1. Search Range
Spanning Trees(1-128) From To

2. Search Options
Bridge Priority(0 = any) State

Search Operation

Spanning Tree Group	State	Bridge Priority	Bridge Hello Time	Bridge Max Age	Bridge Forward Delay	Bridge Aging Time
<u>1</u>	on	32768	2	20	15	300
<u>2</u>	on	32768	2	20	15	300
<u>3</u>	on	32768	2	20	15	300
<u>4</u>	on	32768	2	20	15	300
<u>5</u>	on	32768	2	20	15	300
<u>6</u>	on	32768	2	20	15	300
<u>7</u>	on	32768	2	20	15	300
<u>8</u>	on	32768	2	20	15	300

Select

- d. Enter the Spanning Tree Group number and set the Switch Spanning Tree State to on. To add a VLAN to the Spanning Tree Group, select the VLAN in the VLANs Available list, and click **Add**. VLAN 2 is automatically removed from Spanning Tree Group 1.

Switch Spanning Tree Group Configuration

Spanning Tree Group ID (1-128)	<input type="text" value="2"/>
Switch Spanning Tree State	<input type="button" value="on"/>
Bridge Priority (0-65535)	<input type="text" value="32768"/>
Bridge Hello Time (1-10secs)	<input type="text" value="2"/>
Bridge Max Age (6-40secs)	<input type="text" value="20"/>
Bridge Forward Delay (4-30secs)	<input type="text" value="15"/>

VLANs Available

Vlan ID:Name
 1:Default VLAN
 4095:Mgmt VLAN

VLANs in STG


Vlan ID:Name
 2:VLAN Name

Switch Spanning Tree Port Configuration

Switch Port	Port Priority	Port Path Cost	Port Spanning Tree State
<u>1</u>	128	2	off
<u>2</u>	128	2	off

- e. Scroll down, and click **Submit**.

3. Apply, verify, and save the configuration.



Port Fast Forwarding

Port Fast Forwarding permits a port that participates in Spanning Tree to bypass the Listening and Learning states and enter directly into the Forwarding state. While in the Forwarding state, the port listens to the BPDUs to learn if there is a loop and, if dictated by normal STG behavior (following priorities, etc.), the port transitions into the Blocking state.

This feature permits the switch to interoperate well with Fast Path, a NIC Teaming feature.

Configuring Port Fast Forwarding

Use the following CLI commands to enable Port Fast Forwarding on an external port.

```
>> # /cfg/l2/stp 1/port 20          (Select port 20)
>> Spanning Tree Port 20# fastfwd ena (Enable Port Fast Forwarding)
>> Spanning Tree Port 20# apply      (Make your changes active)
>> Spanning Tree Port 20# save       (Save for restore after reboot)
```

Fast Uplink Convergence

Fast Uplink Convergence enables the switch to quickly recover from the failure of the primary link or trunk group in a Layer 2 network using Spanning Tree Protocol. Normal recovery can take as long as 60 seconds, while the backup link transitions from Blocking to Listening to Learning and then Forwarding states. With Fast Uplink Convergence enabled, the switch immediately places the secondary path into Forwarding state, and sends multicasts of addresses in the forwarding database (FDB) and ARP table over the secondary link so that upstream switches can learn the new path.

Configuration guidelines

When you enable Fast Uplink Convergence, the switch software automatically makes the following configuration changes:

- Increases the bridge priority to 65500 so that it does not become the root switch.
- Increases the cost of all of the external ports by 3000, across all VLANs and Spanning Tree Groups. This ensures that traffic never flows through the switch to get to another switch unless there is no other path.

When you disable Fast Uplink Convergence, the bridge priorities and path cost are set to their default values for all STP groups.

Configuring Fast Uplink Convergence

Use the following CLI commands to enable Fast Uplink Convergence on external ports:

```
>> # /cfg/l2/upfast ena(Enable Fast Uplink convergence)
>> Layer 2# apply(Make your changes active)
>> Layer 2# save(Save for restore after reboot)
```

RSTP and MSTP

Introduction

Rapid Spanning Tree Protocol (IEEE 802.1w) enhances the Spanning Tree Protocol (IEEE 802.1d) to provide rapid convergence on Spanning Tree Group 1. Multiple Spanning Tree Protocol (IEEE 802.1s) extends the Rapid Spanning Tree Protocol to provide both rapid convergence and load balancing in a VLAN environment.

The following topics are discussed in this chapter:

- Rapid Spanning Tree Protocol (RSTP)
- Multiple Spanning Tree Protocol (MSTP)

Rapid Spanning Tree Protocol

Rapid Spanning Tree Protocol (RSTP) provides rapid convergence of the spanning tree and provides for fast reconfiguration critical for networks carrying delay-sensitive traffic such as voice and video. RSTP significantly reduces the time to reconfigure the active topology of the network when changes occur to the physical topology or its configuration parameters. RSTP reduces the bridged-LAN topology to a single Spanning Tree.

For more information about Spanning Tree Protocol, see the “Spanning Tree Protocol” chapter in this guide.

RSTP parameters are configured in Spanning Tree Group 1. STP Groups 2-128 do not apply to RSTP, and must be cleared. There are new STP parameters to support RSTP, and some values to existing parameters are different.

RSTP is compatible with devices that run 802.1d Spanning Tree Protocol. If the switch detects 802.1d BPDUs, it responds with 802.1d-compatible data units. RSTP is not compatible with Per VLAN Spanning Tree (PVST) protocol.

Port state changes

The port state controls the forwarding and learning processes of Spanning Tree. In RSTP, the port state has been consolidated to the following: discarding, learning, and forwarding.

Table 13 RSTP vs. STP port states

Port operational status	STP port state	RSTP port state
Enabled	Blocking	Discarding
Enabled	Listening	Discarding
Enabled	Learning	Learning
Enabled	Forwarding	Forwarding
Disabled	Disabled	Discarding

Port type and link type

Spanning Tree Configuration includes the following parameters to support RSTP and MSTP:

- Edge port
- Link type

Although these parameters are configured for Spanning Tree Groups 1-128 (`/cfg/12/stp y/port x`), they only take effect when RSTP/MSTP is turned on.

Edge port

A port that connects to a server or stub network is called an edge port. Therefore, ports 1-16 should have `edge` enabled. Edge ports can start forwarding as soon as the link is up.

Edge ports do not take part in Spanning Tree, and should not receive BPDUs. If a port with `edge` enabled does receive a BPDU, it begins STP processing until it is re-enabled.

Link type

The link type determines how the port behaves in regard to Rapid Spanning Tree. The link type corresponds to the duplex mode of the port. A full-duplex link is `point-to-point (p2p)`, while a half-duplex link should be configured as `shared`. If you select `auto` as the link type, the port dynamically configures the link type.

RSTP configuration guidelines

This section provides important information about configuring Rapid Spanning Tree Groups:

- When RSTP is turned on, STP parameters apply only to STP Group 1.
- When RSTP is turned on, all VLANs from STP Groups other than STP Group 1 are moved to STP Group 1. The other STP Groups (2-128) are turned off.

RSTP configuration example

This section provides steps to configure Rapid Spanning Tree on the switch, using the Command Line Interface (CLI) or the Browser-based Interface (BBI).

Configuring Rapid Spanning Tree (CLI example)

1. Configure port and VLAN membership on the switch, as described in the “Configuring ports and VLANs (CLI example)” section in the “VLANs” chapter of this guide.
2. Set the Spanning Tree mode to Rapid Spanning Tree.


```
>> /cfg/12/mrst (Select Multiple Spanning Tree menu)
>> Multiple Spanning Tree# mode rstp (Set mode to Rapid Spanning Tree)
>> Multiple Spanning Tree# on (Turn Rapid Spanning Tree on)
```

3. Apply and save the changes.

```
>> # apply (Apply the configuration)
>> # save (Save the configuration)
```

Configuring Rapid Spanning Tree Protocol (BBI example)

1. Configure port and VLAN membership on the switch, as described in the “Configuring ports and VLANs (BBI example)” section in the “VLANs” chapter of this guide.
2. Configure RSTP general parameters.
 - a. Click the **Configure** context button on the Toolbar.
 - b. Open the MSTP/RSTP folder, and select **General**.




- c. Select RSTP mode, and set the MSTP/RSTP state to **ON**.

MSTP/RSTP General Configuration

Region Name	myregion
Revision Level (0-65535)	1
Max. Hop Count (4-60)	20
MSTP/RSTP Mode	RSTP ▼
MSTP/RSTP State	ON ▼

- d. Click **Submit**.

3. Apply, verify, and save the configuration.



Multiple Spanning Tree Protocol

IEEE 802.1s Multiple Spanning Tree extends the IEEE 802.1w Rapid Spanning Tree Protocol through multiple Spanning Tree Groups. MSTP maintains up to 32 spanning-tree instances that correspond to STP Groups 1-32.

In Multiple Spanning Tree Protocol (MSTP), several VLANs can be mapped to each Spanning-Tree instance. Each Spanning-Tree instance is independent of other instances. MSTP allows frames assigned to different VLANs to follow separate paths, each path based on an independent Spanning-Tree instance. This approach provides multiple forwarding paths for data traffic, enabling load balancing, and reducing the number of Spanning-Tree instances required to support a large number of VLANs.

MSTP region

A group of interconnected bridges that share the same attributes is called an MST region. Each bridge within the region must share the following attributes:

- Alphanumeric name
- Revision level
- VLAN-to-STG mapping scheme

MSTP provides rapid reconfiguration, scalability, and control due to the support of regions, and multiple Spanning-Tree instances support within each region.

Common Internal Spanning Tree

The Common Internal Spanning Tree (CIST) provides a common form of Spanning Tree Protocol, with one Spanning Tree instance that can be used throughout the MSTP region. CIST allows the switch to interoperate with legacy equipment, including devices that run IEEE 802.1d (STP).

CIST allows the MSTP region to act as a virtual bridge to other bridges outside of the region, and provides a single Spanning-Tree instance to interact with them.

CIST is the default spanning tree group. When VLANs are removed from STG 1-128, the VLANs automatically become members of the CIST.

CIST port configuration includes Hello time, Edge port status (enable/disable), and Link Type. These parameters do not affect Spanning Tree Groups 1-128. They apply only when the CIST is used.

MSTP configuration guidelines

This section provides important information about configuring Multiple Spanning Tree Groups:

- When you turn on MSTP, the switch automatically moves VLAN 1 to the Common Internal Spanning Tree (CIST).
- Region Name and revision level must be configured. Each bridge in the region must have the same name and revision level.
- The VLAN and STP Group mapping must be the same across all bridges in the region.
- You can move any VLAN to the CIST.
- You can move VLAN 1 into any Spanning Tree Group.

MSTP configuration example

This section provides steps to configure Multiple Spanning Tree Protocol on the switch, using the Command Line Interface (CLI) or the Browser-based Interface (BBI).

Configuring Multiple Spanning Tree Protocol (CLI example)

1. Configure port and VLAN membership on the switch, as described in the “Configuring ports and VLANs (CLI example)” section in the “VLANs” chapter of this guide.
2. Set the mode to Multiple Spanning Tree, and configure MSTP region parameters.

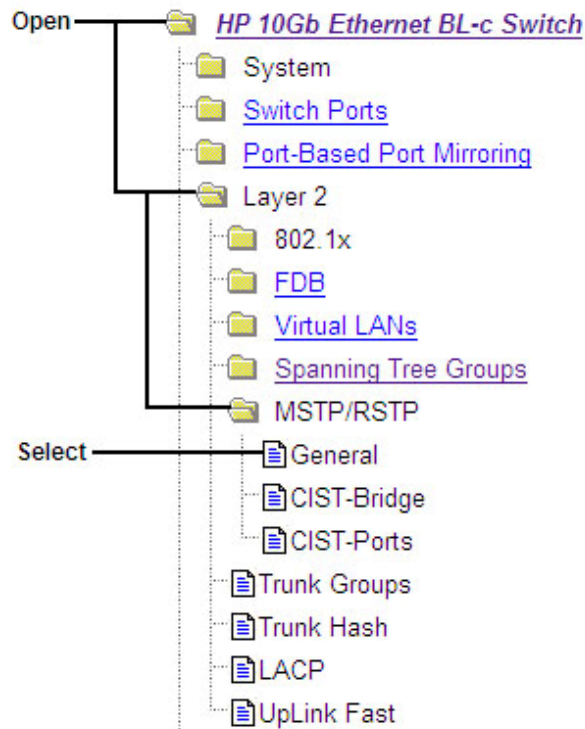
```
>> /cfg/12/ mrst (Select Multiple Spanning Tree menu)
>> Multiple Spanning Tree# mode mstp (Set mode to
Multiple Spanning Trees)
>> Multiple Spanning Tree# on (Turn Multiple Spanning Trees on)
>> Multiple Spanning Tree# name xxxxxx (Define the Region name)
>> Multiple Spanning Tree: rev xx (Define the Region revision level)
```

3. Assign VLANs to Spanning Tree Groups.

```
>> /cfg/12/stp 2 (Select Spanning Tree Group 2)
>> Spanning Tree Group 2# add 2 (Add VLAN 2)
>> Spanning Tree Group 2# apply (Apply the configurations)
```

Configuring Multiple Spanning Tree Protocol (BBI example)

1. Configure port and VLAN membership on the switch, as described in the “Configuring ports and VLANs (BBI example)” section in the “VLANs” chapter of this guide.
2. Configure MSTP general parameters.
 - a. Click the **Configure** context button on the Toolbar.
 - b. Open the MSTP/RSTP folder, and select **General**.



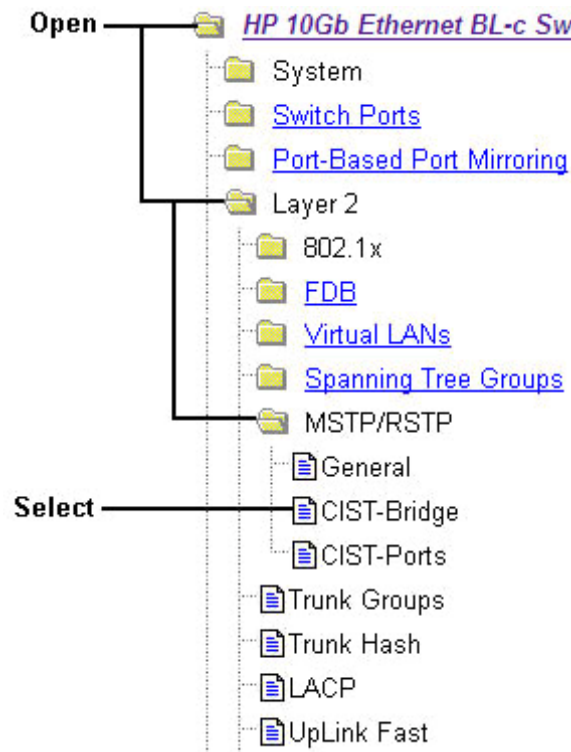
- c. Enter the region name and revision level. Select MSTP mode, and set the MSTP/RSTP state to **ON**.

MSTP/RSTP General Configuration

Region Name	myregion
Revision Level (0-65535)	1
Max. Hop Count (4-60)	20
MSTP/RSTP Mode	MSTP ▼
MSTP/RSTP State	ON ▼

- d. Click **Submit**.

3. Configure Common Internal Spanning Trees (CIST) bridge parameters.
 - a. Open the MSTP/RSTP folder, and select **CIST-Bridge**.



- b. Enter the Bridge Priority, Maximum Age, and Forward Delay values.

Common Internal Spanning Tree Bridge Configuration

Bridge Priority (0-65535)	32768
Max. Age (6-40 secs)	20
Forward Delay (4-30 secs)	15

VLANs Available

Vlan ID:Name

Add>>

Cist VLANs

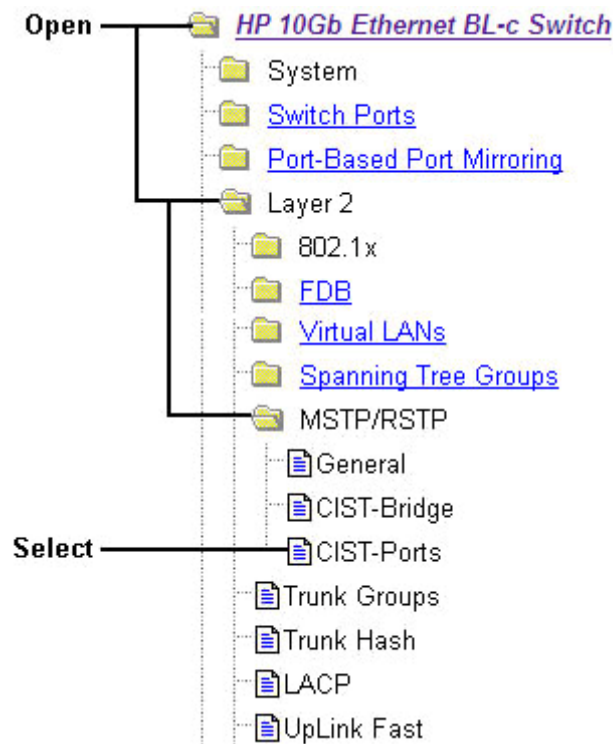
Vlan ID:Name
1:Default VLAN

Submit

- c. Click **Submit**.

4. Configure Common Internal Spanning Tree (CIST) port parameters.

- a. Open the MSTP/RSTP folder, and select **CIST-Ports**.



- b. Click a port number to select it.

Ports Common Internal Spanning Tree Configuration

CIST Port	Priority	Port Path Cost	Link Type	Edge Port State	Port STP State
<u>1</u>	128	2000	auto	enabled	OFF
<u>2</u>	128	2000	auto	enabled	OFF
<u>3</u>	128	2000	auto	enabled	OFF
<u>4</u>	128	2000	auto	enabled	OFF
<u>5</u>	128	2000	auto	enabled	OFF
<u>6</u>	128	2000	auto	enabled	OFF
<u>7</u>	128	2000	auto	enabled	OFF
<u>8</u>	128	2000	auto	enabled	OFF
<u>9</u>	128	2000	auto	enabled	OFF
<u>10</u>	128	2000	auto	enabled	OFF


An arrow labeled 'Select' points to the first row of the table, specifically to the port number '1'.

- c. Enter the Port Priority, Path Cost, and select the Link Type. Set the CIST Port State to **ON**.

Common Internal Spanning Tree Port 1 Configuration

Port Priority (0-240)	<input type="text" value="128"/>
Path Cost (1-200000000, 0 for auto)	<input type="text" value="2000"/>
Link Type	Auto <input type="button" value="v"/>
Enable/Disable Edge	Enabled <input type="button" value="v"/>
Port STP State	ON <input type="button" value="v"/>
Hello Time (1-10 secs)	<input type="text" value="2"/>

- d. Click **Submit**.
5. Apply, verify, and save the configuration.



Quality of Service

Introduction

Quality of Service features allow you to allocate network resources to mission-critical applications at the expense of applications that are less sensitive to such factors as time delays or network congestion. You can configure your network to prioritize specific types of traffic, ensuring that each type receives the appropriate Quality of Service (QoS) level.

The following topics are discussed in this section:

- Quality of Service Overview
- Using ACL Filters
- Using DSCP Values to Provide QoS
- Using 802.1p Priorities to Provide QoS
- Queuing and Scheduling


Overview

QoS helps you allocate guaranteed bandwidth to the critical applications, and limit bandwidth for less critical applications. Applications such as video and voice must have a certain amount of bandwidth to work correctly; using QoS, you can provide that bandwidth when necessary. Traffic for applications that are sensitive to timing out or cannot tolerate delay can be assigned to a high-priority queue.

By assigning QoS levels to traffic flows on your network, you can ensure that network resources are allocated where they are needed most. QoS features allow you to prioritize network traffic, thereby providing better service for selected applications.

The following figure shows the basic QoS model used by the HP 10GbE switch.

Figure 11 QoS model



The switch uses the Differentiated Services (*DiffServ*) architecture to provide QoS functions. *DiffServ* is described in IETF RFCs 2474 and 2475.

With *DiffServ*, you can establish policies to direct traffic. A policy is a traffic-controlling mechanism that monitors the characteristics of the traffic, (for example, its source, destination, and protocol) and performs a controlling action on the traffic when certain characteristics are matched.

The switch can classify traffic by reading the IEEE 802.1p priority value, or by using filters to match specific criteria. When network traffic attributes match those specified in a traffic pattern, the policy instructs the switch to perform specified actions on each packet that passes through it. The packets are assigned to different Class of Service (COS) queues and scheduled for transmission.

The basic HP 10GbE switch QoS model works as follows:

- Classify traffic:
 - Read 802.1p Priority.
 - Match ACL filter parameters.
- Meter traffic:
 - Define bandwidth and burst parameters.
 - Select actions to perform on in-profile and out-of-profile traffic.
- Perform actions:
 - Drop packets.
 - Pass packets.
 - Mark DSCP or 802.1p Priority.
 - Set COS queue (with or without re-marking).
- Queue and schedule traffic:
 - Place packets in a COS queue (configurable as either 2 or 8 queues).
 - Schedule transmission based on the COS queue weight.

Using ACL filters

Access Control Lists are filters that allow you to classify and segment traffic, so you can provide different levels of service to different traffic types. Each filter defines the conditions that must match for inclusion in the filter, and also the actions that are performed when a match is made.

Summary of packet classifiers

The HP 10GbE switch allows you to classify packets based on various parameters, such as:

- Ethernet
 - Source MAC address/mask
 - Destination MAC address/mask
 - VLAN number/mask
 - Ethernet type
 - Ethernet Priority, which is the IEEE 802.1p Priority
- IPv4
 - Source IP address/mask
 - Destination IP address/mask
 - Type of Service value
 - IP protocol number: The protocol number or name as shown in the following table:

Table 14 Well-known protocol types

Number	Protocol Name
1	icmp
2	igmp
6	tcp
17	udp

Table 14 Well-known protocol types

Number	Protocol Name
89	ospf
112	vrrp

- TCP/UDP
 - TCP/UDP application source port, as shown in the table titled “Well-Known Application Ports”
 - TCP/UDP application destination port, as shown in the table titled “Well-Known Application Ports”
 - TCP/UDP flag value, as shown in the table titled “Well-Known TCP Flag Values”

Table 15 Well-known application ports

Number	TCP/UDP Application	Number	TCP/UDP Application	Number	TCP/UDP Application
20	ftp-data	79	finger	179	irc
21	ftp	80	http	194	imap3
22	ssh	109	pop2	220	ldap
23	telnet	110	pop3	389	https
25	smtp	111	sunrpc	443	rip
37	time	119	nntp	520	rtsp
42	name	123	ntp	554	Radius
43	whois	143	imap	1645;1812	Radius Accounting
53	domain	144	news	1813	hsrp
69	tftp	161	snmp	1985	
70	gopher	162	snmptrap		

Table 16 Well-known TCP flag values

Flag	Value
URG	0x0020
ACK	0x0010
PSH	0x0008
RST	0x0004
SYN	0x0002
FIN	0x0001

- Packet Format
 - Ethernet format (eth2, SNAP, LLC)
 - Ethernet tagging format
- Egress port packets

Note that the egress port ACL will not match a broadcast, multicast, unknown unicast, or Layer 3 packet. The egress port ACL will not match packets if the destination port is a trunk member.

Summary of ACL actions

Actions determine how the traffic is treated. The HP 10GbE switch QoS actions include the following:

- Pass or Drop
- Re-mark a new DiffServ Code Point (DSCP)
- Re-mark the 802.1p field
- Set the COS queue

Understanding ACL precedence

Each ACL has a unique precedence level, based on its number. When an incoming packet matches the highest precedence ACL, the ACL's configured action takes place. The other assigned ACLs also are considered, in order of precedence.

ACLs are divided into Precedence Groups, as shown in the following table.

Precedence Group	ACLs	Precedence Level
Precedence Group 1	ACL 1 – ACL 128	Low
Precedence Group 2	ACL 129 – ACL 256	
Precedence Group 3	ACL 257 – ACL 384	High

NOTE: Precedence Groups are not related to ACL Groups.

Each Precedence Group has its own precedence level, such that Precedence Group 2 has a higher precedence level than Precedence Group 1. Within each Precedence Group, higher-numbered ACLs receive higher precedence, so that the lowest-numbered ACL has the lowest precedence level, and the highest-numbered ACL has the highest precedence level. However, the other ACLs within the Precedence Group have an unspecified precedence level, as follows:

ACL 1 = lowest precedence level within Precedence Group 1
 ACL 2 = unspecified precedence level within Precedence Group 1
 ACL 3 = unspecified precedence level within Precedence Group 1
 ...
 ACL 126 = unspecified precedence level within Precedence Group 1
 ACL 127 = unspecified precedence level within Precedence Group 1
 ACL 128 = highest precedence level within Precedence Group 1

Using ACL Groups

Access Control Lists (ACLs) allow you to classify packets according to a particular content in the packet header, such as the source address, destination address, source port number, destination port number, and others. Packet classifiers identify flows for more processing.

You can define a traffic profile by compiling a number of ACLs into an ACL Group, and assigning the ACL Group to a port.

ACL Groups are assigned and enabled on a per-port basis. Each ACL can be used by itself or in combination with other ACLs or ACL Groups on a given switch port.

ACLs can be grouped in the following manner:

- Access Control Lists

Access Control Lists (ACLs) allow you to classify packets according to a particular content in the packet header, such as the source address, destination address, source port number, destination port number, and others. Packet classifiers identify flows for more processing.

The HP 10GbE switch supports up to 384 ACLs. Each ACL defines one filter rule. Each filter rule is a collection of matching criteria, and can include an action (permit or deny the packet). For example:

```
ACL 200:
VLAN = 1
SIP = 10.10.10.1 (255.255.255.0)
Action = permit
```

- Access Control Groups

An Access Control Group (ACL Group) is a collection of ACLs. For example:

```
ACL Group 1

ACL 382:
VLAN = 1
SIP = 10.10.10.1 (255.255.255.0)
Action = permit

ACL 383:
VLAN = 2
SIP = 10.10.10.2 (255.255.255.0)
Action = deny

ACL 384:
PRI = 7
DIP = 10.10.10.3 (255.255.0.0)
Action = permit
```

In the example above, each ACL defines a filter rule. ACL 383 has a higher precedence than ACL 382, based on its number.

Use ACL Groups to create a traffic profile by gathering ACLs into an ACL Group, and assigning the ACL Group to a port. The HP 10GbE switch supports up to 384 ACL Groups. Each ACL group supports up to 96 ACLs.

ACL Metering and Re-marking

You can define a profile for the aggregate traffic flowing through the HP 10GbE switch, by configuring a QoS meter (if desired), and assigning ACL Groups to ports. When you add ACL Groups to a port, make sure they are ordered correctly in terms of precedence.

For example, consider two ACL Groups, ACL Group 1 and ACL Group 2. Each contains three levels of precedence. If you add ACL Group 1 to a port, then add ACL Group 2 to the port, the port's ACL filters contain a total of six precedence levels. ACL Group 1 has precedence over ACL Group 2.

Each port supports up to seven precedence levels.

Actions taken by an ACL are called *In-Profile* actions. You can configure additional In-Profile and Out-of-Profile actions on a port. Data traffic can be metered, and re-marked to ensure that the traffic flow provides certain levels of service in terms of bandwidth for different types of network traffic.

Metering

QoS metering provides different levels of service to data streams through user-configurable parameters. A meter is used to measure the traffic stream against a traffic profile, which you create. Thus, creating meters yields In-Profile and Out-of-Profile traffic for each ACL, as follows:

- In-Profile—If there is no meter configured or if the packet conforms to the meter, the packet is classified as In-Profile.
- Out-of-Profile—If a meter is configured and the packet does not conform to the meter (exceeds the committed rate or maximum burst rate of the meter), the packet is classified as Out-of-Profile.

Using meters, you set a Committed Rate in Kb/s (1024 bits per second in each Kb/s). All traffic within this Committed Rate is In-Profile. Additionally, you set a Maximum Burst Size that specifies an allowed data burst larger than the Committed Rate for a brief period. These parameters define the In-Profile traffic.

Meters keep the sorted packets within certain parameters. You can configure a meter on an ACL, and perform actions on metered traffic, such as packet re-marking.

Re-marking

Re-marking allows for the treatment of packets to be reset based on new network specifications or desired levels of service. You can configure the ACL to re-mark a packet as follows:

- Change the DSCP value of a packet, used to specify the service level traffic should receive.
- Change the 802.1p priority of a packet.

Viewing ACL statistics

ACL statistics display how many packets hit (matched) each ACL. Up to 64 statistic counters can be displayed for each ACL Precedence Group. Use ACL statistics to check filter performance, and debug the ACL filters.

You must enable statistics (`cfg/acl/acl x/stats ena`) for each ACL that you want to monitor.

ACL configuration examples

Configure Access Control Lists (CLI example)

The following configuration examples illustrate how to use Access Control Lists (ACLs) to block traffic. These basic configurations illustrate common principles of ACL filtering.

NOTE: Each ACL filters traffic that ingresses on the port to which the ACL is added. The egrport classifier filters traffic that ingresses the port to which the ACL is added, and then egresses the port specified by egrport. In most common configurations, egrport is not used.

- **Example 1** Use this configuration to block traffic to a specific host.

```
>> Main# /cfg/acl/acl 255                                (Define ACL 255)
>> ACL 255# ipv4/dip 100.10.1.116 255.255.255.255
>> Filtering IPv4# ..
>> ACL 255# action deny
>> ACL 255# /cfg/port 20/aclqos                         (Add ACL to port 20)
>> Port 20 ACL# add acl 255
>> Port 20 ACL# apply
>> Port 20 ACL# save
```

In this example, all traffic that ingresses on port 20 is denied if it is destined for the host at IP address 100.10.1.116.

- **Example 2** Use this configuration to block traffic from a network destined for a specific host address.

```
>> Main# /cfg/acl/acl 256                                (Define ACL 256)
>> ACL 256# ipv4/sip 100.10.1.0 255.255.255.0
>> ACL 256# ipv4/dip 200.20.1.116 255.255.255.255
>> Filtering IPv4# ..
>> ACL 256# action deny
>> ACL 256# /cfg/port 20/aclqos                         (Add ACL to port 20)
>> Port 20 ACL# add acl 256
>> Port 20 ACL# apply
>> Port 20 ACL# save
```

In this example, all traffic that ingresses on port 20 with source IP from the class 100.10.1.0/24 and destination IP 200.20.1.116 is denied.

- **Example 3** Use this configuration to block traffic from a source that is destined for a specific egress port.

```


>> Main# /cfg/acl/acl 1                                (Define ACL 1)
>> ACL 1# ethernet/smac 00:21:00:00:00:00 ff:ff:ff:ff:ff:ff
>> Filtering Ethernet# ..
>> ACL 1# action deny
>> ACL 1# stats e
>> ACL 1# /cfg/acl/acl 257                            (Define ACL 257)
>> ACL 257# egrport 21
>> ACL 257# action deny
>> ACL 257# stats e
>> ACL 257# /cfg/port 20/aclqos
>> Port 20 ACL# add acl 1                             (Add ACL 1 to port 20)
>> Port 20 ACL# add acl 257                          (Add ACL 257 to port 20)
>> Port 20 ACL# apply
>> Port 20 ACL# save

```

In this example, all traffic (Layer 2 known unicast) that ingresses on port 20 from source MAC 00:21:00:00:00:00 and is destined for port 21 is denied.

Configure Access Control Lists and Groups (BBI example 1)

1. Configure Access Control Lists (ACLs).
 - a. Click the **Configure** context button on the Toolbar.
 - b. Open the Access Control Lists folder, and select **Add ACL**.




- c. Configure the ACL parameters. Set the Filter Action to **Deny**, the Ethernet Type to **IPv4**, and the Destination IP Address to **100.10.1.116**.

Access Control List

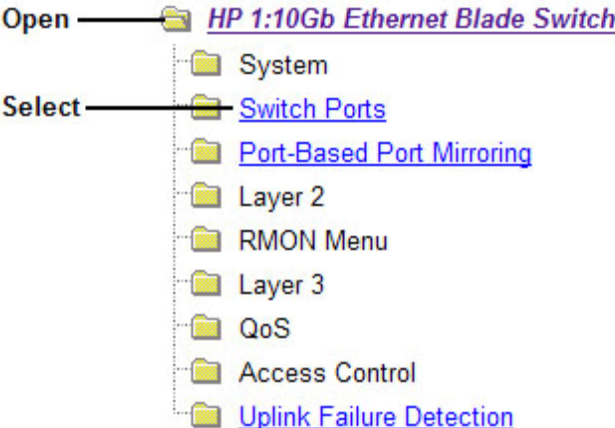
ACL Id (1 - 384)	255	Group Id: 0
Filter Action	Deny	Set priority value none
Ethernet Packet Format	Disabled	
Tagging Packet Format	Disabled	
Source MAC Address	00:00:00:00:00:00	Mask ff:ff:ff:ff:ff:ff
Destination MAC Address	00:00:00:00:00:00	Mask ff:ff:ff:ff:ff:ff
Ethernet Type	None	Value (0600-ffff) 600
VLAN Id (1-4095)	1	Mask (0-fff) fff Disabled
802.1p Priority	None	
Type of Service (0-255)	0	Disabled
Protocol (0-255)	0	Disabled
Source IP Address	0.0.0.0	Mask 255.255.255.255
Destination IP Address	100.10.1.116	Mask 255.255.255.255
TCP/UDP Src Port (1-65535)	1	Mask (0-ffff) ffff Disabled
TCP/UDP Dst Port (1-65535)	1	Mask (0-ffff) ffff Disabled
TCP Flags	<input type="checkbox"/> FIN <input type="checkbox"/> SYN <input type="checkbox"/> RST <input type="checkbox"/> PSH <input type="checkbox"/> ACK <input type="checkbox"/> URG	Mask(0-3f) 3f Disabled
Statistics	Disabled	
Egress port	None	

- d. Click **Submit**.

2. Apply, verify, and save the configuration.



- 3. Add ACL 1 to port 1.
 - a. Click the **Configure** context button on the Toolbar.
 - b. Select **Switch Ports** (click the underlined text, not the folder).



- c. Select a port.

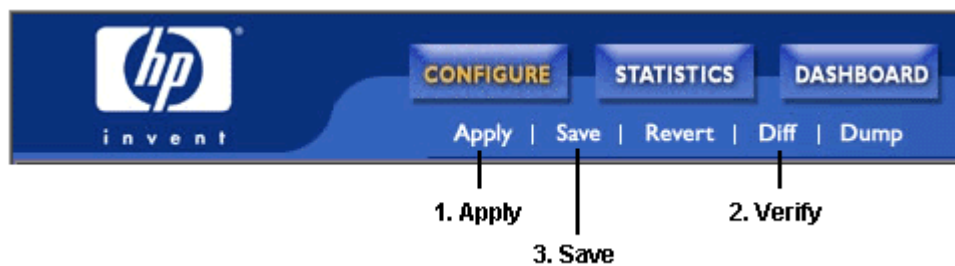
Switch Port	State	VLAN Tagging	Default PVID	PVID tagging	Multicast Threshold	Broadcast Threshold	Destination Lookup Fail Threshold	802.1p Priority
<u>1</u>	enabled	disabled	1	enabled	disabled	disabled	disabled	0
<u>2</u>	enabled	disabled	2	enabled	disabled	disabled	disabled	0
<u>3</u>	enabled	disabled	1	enabled	disabled	disabled	disabled	0
<u>4</u>	enabled	disabled	1	enabled	disabled	disabled	disabled	0
<u>5</u>	enabled	disabled	1	enabled	disabled	disabled	disabled	0
<u>6</u>	enabled	disabled	1	enabled	disabled	disabled	disabled	0
<u>7</u>	enabled	disabled	1	enabled	disabled	disabled	disabled	0
<u>8</u>	enabled	disabled	1	enabled	disabled	disabled	disabled	0
<u>9</u>	enabled	disabled	1	enabled	disabled	disabled	disabled	0
<u>10</u>	enabled	disabled	1	enabled	disabled	disabled	disabled	0

- d. Add the ACL to the port.

The screenshot displays two configuration panels. The top panel, titled "Switch Port 1 Configuration", shows the "Switch Port State" set to "Enabled". The bottom panel, titled "ACL Configuration", is divided into four sections: "ACLs Available", "ACLs Selected", "ACL Groups Available", and "ACL Groups Selected". In the "ACLs Selected" section, the "ACL ID" dropdown menu is open, showing the value "255". In the "ACL Groups Selected" section, the "ACL Group ID" dropdown menu is also open. Between the "ACLs Available" and "ACLs Selected" sections, there are "Add >>" and "<< Remove" buttons. Similarly, between the "ACL Groups Available" and "ACL Groups Selected" sections, there are "Add >>" and "<< Remove" buttons. A "Submit" button is located at the bottom center of the "ACL Configuration" panel.

- e. Click **Submit**.

4. Apply, verify, and save the configuration.



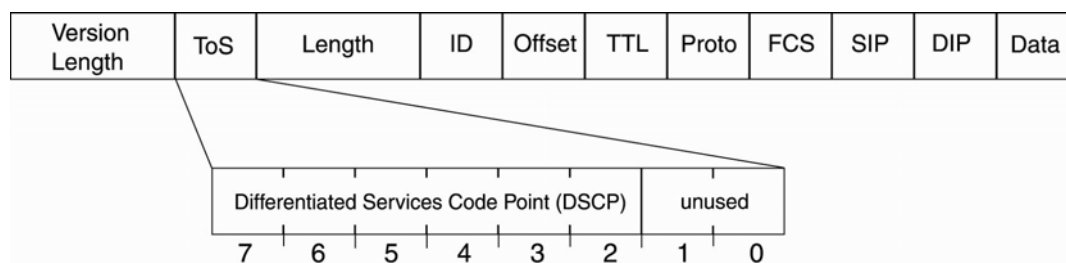
Using DSCP values to provide QoS

The six most significant bits in the TOS byte of the IP header are defined as DiffServ Code Points (DSCP). Packets are marked with a certain value depending on the type of treatment the packet must receive in the network device. DSCP is a measure of the Quality of Service (QoS) level of the packet.

Differentiated Services concepts

To differentiate between traffic flows, packets can be classified by their DSCP value. The Differentiated Services (DS) field in the IP header is an octet, and the first six bits, called the DS Code Point (DSCP), can provide QoS functions. Each packet carries its own QoS state in the DSCP. There are 64 possible DSCP values (0-63).

Figure 12 Layer 3 IPv4 packet



The switch can perform the following actions to the DSCP:

- Read the DSCP value of ingress packets
- Re-mark the DSCP value to a new value
- Map the DSCP value to an 802.1p priority

Once the DSCP value is marked, the switch can use it to direct traffic prioritization.

Per Hop Behavior

The DSCP value determines the Per Hop Behavior (PHB) of each packet. The PHB is the forwarding treatment given to packets at each hop. QoS policies are built by applying a set of rules to packets, based on the DSCP value, as they hop through the network.

The HP 10GbE switch default settings are based on the following standard PHBs, as defined in the IEEE standards:

- Expedited Forwarding (EF)—This PHB has the highest egress priority and lowest drop precedence level. EF traffic is forwarded ahead of all other traffic. EF PHB is described in RFC 2598.
- Assured Forwarding (AF)—This PHB contains four service levels, each with a different drop precedence, as shown below. Routers use drop precedence to determine which packets to discard last when the network becomes congested. AF PHB is described in RFC 2597.

Table 17 Assured forwarding drop-down precedence

Drop Precedence	Class 1	Class 2	Class 3	Class 4
Low	AF11 (DSCP 10)	AF21 (DSCP 18)	AF31 (DSCP 26)	AF41 (DSCP 34)
Medium	AF12 (DSCP 12)	AF22 (DSCP 20)	AF32 (DSCP 28)	AF42 (DSCP 36)
High	AF13 (DSCP 14)	AF23 (DSCP 22)	AF33 (DSCP 30)	AF43 (DSCP 38)

- Class Selector (CS)—This PHB has eight priority classes, with CS7 representing the highest priority, and CS0 representing the lowest priority, as shown below. CS PHB is described in RFC 2474.

Table 18 Class selector priority classes

Priority	Class Selector	DSCP
Highest	CS7	56
	CS6	48
	CS5	40
	CS4	32
	CS3	24
	CS2	16
	CS1	8
Lowest	CS0	0

QoS levels

The following table shows the default service levels provided by the switch, listed from highest to lowest importance:

Table 19 Default QoS service levels

Service Level	Default PHB	802.1p Priority
Critical	CS7	7
Network Control	CS6	6
Premium	EF, CS5	5
Platinum	AF41, AF42, AF43, CS4	4
Gold	AF31, AF32, AF33, CS3	3
Silver	AF21, AF22, AF23, CS2	2
Bronze	AF11, AF12, AF13, CS1	1


Using 802.1p priorities to provide QoS

The HP 10GbE switch software provides Quality of Service functions based on the priority bits in a packet's VLAN header. (The priority bits are defined by the 802.1p standard within the IEEE 802.1q VLAN header.) The 802.1p bits, if present in the packet, specify the priority given to packets during forwarding. Packets with a numerically higher (non-zero) priority are given forwarding preference over packets with lower priority.

Packets with a priority mapped to a higher Class of Service (COS) and COS queue (COSq) weight are given forwarding preference over packets with priority mapped to a lower COS and COSq weight. The scheduling scheme is Weight Round Robin (WRR), with user-configurable weight from 0 to 15 for a COSq. The switch can be configured with either two or eight output Class of Service queues (COSq).

The IEEE 802.1p standard uses eight levels of priority (0-7). Priority 7 is assigned to highest priority network traffic, such as OSPF or RIP routing table updates, priorities 5-6 are assigned to delay-sensitive applications such as voice and video, and lower priorities are assigned to standard applications. A value of 0 (zero) indicates a best effort traffic prioritization, and this is the default when traffic priority has not been configured on your network. The switch can filter packets based on the 802.1p values, and it can assign or overwrite the 802.1p value in the packet.

Figure 13 Layer 2 802.1q/802.1p VLAN tagged packet



Ingress packets receive a priority value, as follows:

- Tagged packets—the switch reads the 802.1p priority in the VLAN tag.
- Untagged packets—the switch tags the packet and assigns an 802.1p priority, based on the port's default priority (`/cfg/port x/8021ppri`).

Egress packets are placed in a COS queue based on the priority value, and scheduled for transmission based on the scheduling weight of the COS queue.

Use the `/cfg/qos/8021p/cur` command to display the mapping between 802.1p values, Class of Service queues (COSq), and COSq scheduling weights.

```
>> 802.1p# cur
Current priority to COS queue configuration:
Number of COSq: 2
Priority  COSq  Weight
-----  -
0         0         1
1         0         1
2         0         1
3         0         1
4         1         2
5         1         2
6         1         2
7         1         2
```

802.1p configuration (CLI example)

1. Configure a port's default 802.1 priority.

```
>> Main# cfg/port 20      (Select port)
>> Port 20# 8021ppri     (Set port's default 802.1p priority)
Current 802.1p priority: 0
Enter new 802.1p priority [0-7]: 1

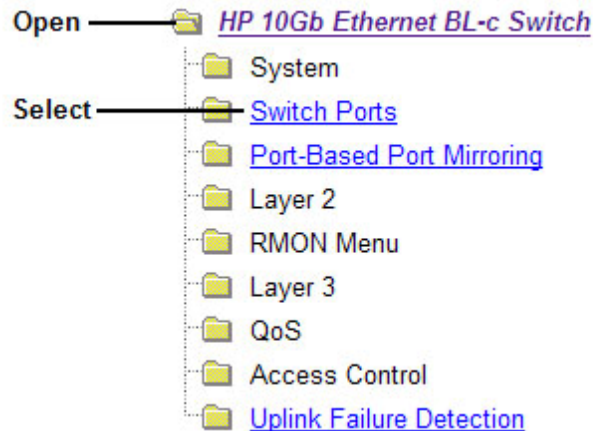
>> Port 20# apply
```

2. Map the 802.1p priority value to a COS queue and set the COS queue scheduling weight.


```
>> Main# cfg/qos/8021p   (Select 802.1p menu)
>> 802.1p# priq         (Set COS queue assignments)
Enter priority [0-7]: 1
Current COS queue (for priority 1): 0
Enter new COS queue (for priority 1) [0-1]: 1
>> 802.1p# qweight      (Set COS queue weights)
Enter COS queue [0-2]: 1
Current weight (for COS queue 1): 0
Enter new weight (for COS queue 1) [0-15]: 1
>> 802.1p# apply
```

802.1p configuration (BBI example)

1. Configure a port's default 802.1p priority.
 - a. Click the **Configure** context button on the Toolbar.
 - b. Select **Switch Ports** (click the underlined text, not the folder).



- c. Select a port.



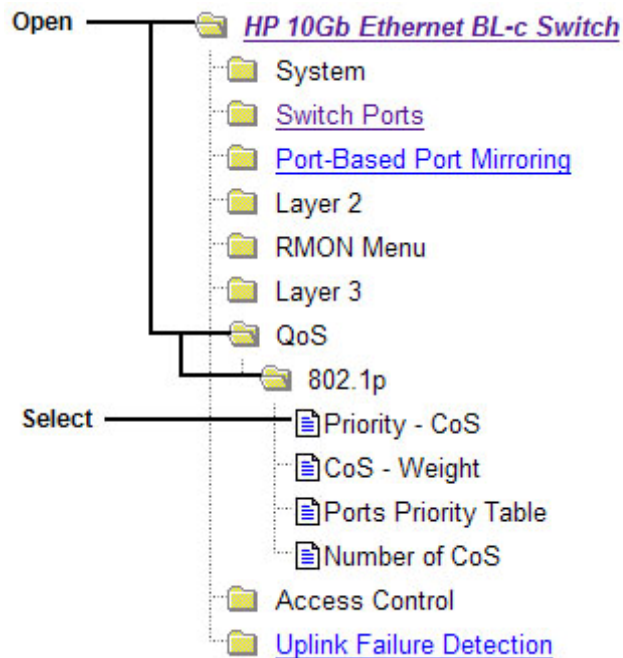
Switch Port	State	VLAN Tagging	Default PVID	PVID tagging	Multicast Threshold	Broadcast Threshold	Destination Lookup Fail Threshold	802.1p Priority
<u>1</u>	enabled	disabled	1	enabled	disabled	disabled	disabled	0
<u>2</u>	enabled	disabled	2	enabled	disabled	disabled	disabled	0
<u>3</u>	enabled	disabled	1	enabled	disabled	disabled	disabled	0
<u>4</u>	enabled	disabled	1	enabled	disabled	disabled	disabled	0
<u>5</u>	enabled	disabled	1	enabled	disabled	disabled	disabled	0
<u>6</u>	enabled	disabled	1	enabled	disabled	disabled	disabled	0
<u>7</u>	enabled	disabled	1	enabled	disabled	disabled	disabled	0
<u>8</u>	enabled	disabled	1	enabled	disabled	disabled	disabled	0
<u>9</u>	enabled	disabled	1	enabled	disabled	disabled	disabled	0
<u>10</u>	enabled	disabled	1	enabled	disabled	disabled	disabled	0

- d. Set the 802.1p priority value.

Switch Port 1 Configuration	
Switch Port State	Enabled ▾
RMON Instrumentation	Disabled ▾
VLAN Tagging	Disabled ▾
PVID Tagging	Enabled ▾
Port STP	Off ▾
Default Port VLAN ID (1 - 4095)	1
Flow Control	both Rx/Tx ▾
Autonegotiation	Off
Speed	10000
Duplex Mode	Full
Enable/Disable sending Link UP/Down Trap	Enabled ▾
Port Name	Downlink1
Multicast Threshold	Disabled ▾
Multicast Threshold Rate (0-262143)	0
Broadcast Threshold	Disabled ▾
Broadcast Threshold Rate (0-262143)	0
Destination Lookup Fail Threshold	Disabled ▾
Destination Lookup Fail Threshold Rate (0-262143)	0
802.1p Port Priority (0-7)	0

- e. Click **Submit**.

2. Map the 802.1p priority value to a COS queue.
 - a. Click the **Configure** context button on the Toolbar.
 - b. Open the 802.1p folder, and select **Priority - CoS**.



- c. Select an 802.1p priority value.

Priority CoS Configuration Table

Priority	CoS
<u>0</u>	0
<u>1</u>	0
<u>2</u>	0
<u>3</u>	0
<u>4</u>	1
<u>5</u>	1
<u>6</u>	1
<u>7</u>	1

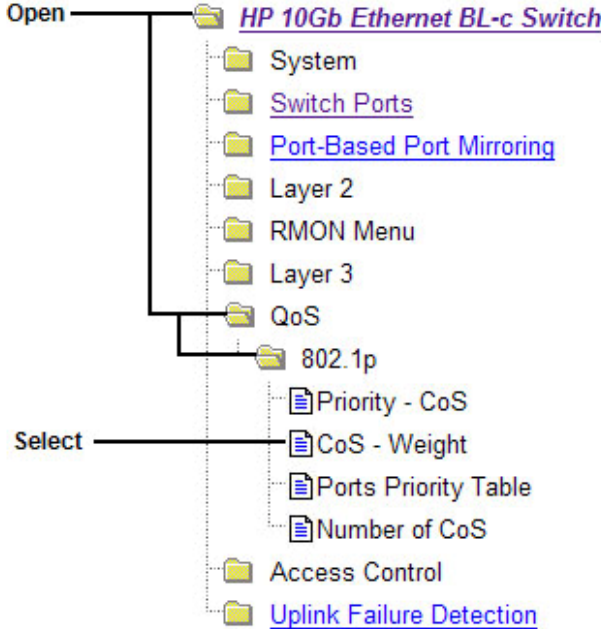
- d. Select a Class of Service queue (CoSQ) to correlate with the 802.1p priority value.

CoSQ For Priority 0 Configuration

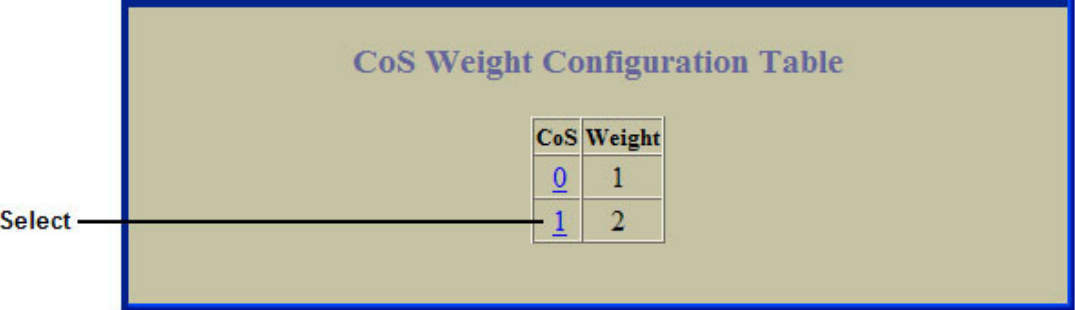
CoSQ (0-1)

- e. Click **Submit**.


- 3. Set the COS queue scheduling weight.
 - a. Click the **Configure** context button on the Toolbar.
 - b. Open the 802.1p folder, and select **CoS - Weight**.




- c. Select a Class of Service queue (CoS).



- d. Enter a value for the weight of the Class of Service queue.



- e. Click **Submit**.
4. Apply, verify, and save the configuration.



Queuing and scheduling

The switch can be configured with either two or eight output Class of Service queues (COSq), into which each packet is placed. Each packet's 802.1p priority determines its COSq, except when an ACL action sets the COSq of the packet.

Each COS queue uses Weighted Round Robin (WRR) scheduling, with user-configurable weight from 0 to 15. The weight of 0 (zero) indicates strict priority, which might starve the low priority queues.

You can configure the following attributes for COS queues:

- Map 802.1p priority value to a COS queue.
- Define the scheduling weight of each COS queue.

Use the 802.1p menu (`/cfg/qos/8021p`) to configure Class of Service queues.

Basic IP routing

This chapter provides configuration background and examples for using the HP 10GbE switch to perform IP routing functions. The following topics are addressed in this chapter:

- IP Routing Benefits
- Routing Between IP Subnets
- Example of Subnet Routing
- Defining IP Address Ranges for the Local Route Cache
- Dynamic Host Configuration Protocol

IP routing benefits

The switch uses a combination of configurable IP switch interfaces and IP routing options. The switch IP routing capabilities provide the following benefits:

- Connects the server IP subnets to the rest of the backbone network.
- Provides another means to invisibly introduce Jumbo frame technology into the server-switched network by automatically fragmenting UDP Jumbo frames when routing to non-Jumbo frame VLANs or subnets.
- Provides the ability to route IP traffic between multiple Virtual Local Area Networks (VLANs) configured on the switch.


Routing between IP subnets

The physical layout of most corporate networks has evolved over time. Classic hub/router topologies have given way to faster switched topologies, particularly now that switches are increasingly intelligent. HP 10GbE switches are intelligent and fast enough to perform routing functions on a par with wire speed Layer 2 switching.

The combination of faster routing and switching in a single device provides another service—it allows you to build versatile topologies that account for legacy configurations.

For example, consider the following topology migration:

Figure 14 Router legacy network



In this example, a corporate campus has migrated from a router-centric topology to a faster, more powerful, switch-based topology. As is often the case, the legacy of network growth and redesign has left the system with a mix of illogically distributed subnets.

This is a situation that switching alone cannot cure. Instead, the router is flooded with cross-subnet communication. This compromises efficiency in two ways:


- Routers can be slower than switches. The cross-subnet side trip from the switch to the router and back again adds two hops for the data, slowing throughput considerably.
- Traffic to the router increases, increasing congestion.

Even if every end-station could be moved to better logical subnets (a daunting task), competition for access to common server pools on different subnets still burdens the routers.

This problem is solved by using HP 10GbE switch with built-in IP routing capabilities. Cross-subnet LAN traffic can now be routed within the switches with wire speed Layer 2 switching performance. This not only eases the load on the router but saves the network administrators from reconfiguring each and every end-station with new IP addresses.

Take a closer look at the HP 10GbE switch in the following configuration example:

Figure 15 Switch-based routing topology



The switch connects the Gigabit Ethernet and Fast Ethernet trunks from various switched subnets throughout one building. Common servers are placed on another subnet attached to the switch. Primary and backup routers are attached to the switch on yet another subnet.

Without Layer 3 IP routing on the switch, cross-subnet communication is relayed to the default gateway (in this case, the router) for the next level of routing intelligence. The router fills in the necessary address information and sends the data back to the switch, which then relays the packet to the proper destination subnet using Layer 2 switching.

With Layer 3 IP routing in place on the switch, routing between different IP subnets can be accomplished entirely within the switch. This leaves the routers free to handle inbound and outbound traffic for this group of subnets.

Example of subnet routing

Prior to configuring, you must be connected to the switch Command Line Interface (CLI) as the administrator.

NOTE: For details about accessing and using any of the menu commands described in this example, see the *HP 10Gb Ethernet BL-c Switch Command Reference*.

1. Assign an IP address (or document the existing one) for each router and client workstation.
2. In the example topology, the following IP addresses are used:

Table 20 Subnet routing example: IP address assignments

Subnet	Devices	IP Addresses
1	Primary and Secondary Default Routers	205.21.17.1 and 205.21.17.2
2	First Floor Client Workstations	100.20.10.2-254
3	Second Floor Client Workstations	131.15.15.2-254
4	Common Servers	206.30.15.2-254

3. Assign an IP interface for each subnet attached to the switch.
4. Since there are four IP subnets connected to the switch, four IP interfaces are needed:

Table 21 Subnet routing example: IP interface assignments

Interface	Devices	IP Interface Address
IF 1	Primary and Secondary Default Routers	205.21.17.3
IF 2	First Floor Client Workstations	100.20.10.1
IF 3	Second Floor Client Workstations	131.15.15.1
IF 4	Common Servers	206.30.15.1

5. IP interfaces are configured using the following commands at the CLI:

```
>> # /cfg/l3/if 1 (Select IP interface 1)
>> IP Interface 1# addr 205.21.17.3 (Assign IP address)
>> IP Interface 1# ena (Enable IP interface 1)
>> IP Interface 1# ../if 2 (Select IP interface 2)
>> IP Interface 2# addr 100.20.10.1 (Assign IP address)
>> IP Interface 2# ena (Enable IP interface 2)
>> IP Interface 2# ../if 3 (Select IP interface 3)
>> IP Interface 3# addr 131.15.15.1 (Assign IP address)
>> IP Interface 3# ena (Enable IP interface 3)
>> IP Interface 3# ../if 4 (Select IP interface 4)
>> IP Interface 4# addr 206.30.15.1 (Assign IP address)
>> IP Interface 4# ena (Enable IP interface 5)
```

6. Set each server and workstation's default gateway to the appropriate switch IP interface (the one in the same subnet as the server or workstation).
7. Configure the default gateways to the routers' addresses.

8. Configuring the default gateways allows the switch to send outbound traffic to the routers:

```
>> IP Interface 5# ../gw 1          (Select primary default gateway)
>> Default gateway 1# addr 205.21.17.1 (Assign IP address)
>> Default gateway 1# ena          (Enable primary default gateway)
>> Default gateway 1# ../gw 2      (Select secondary default gateway)
>> Default gateway 2# addr 205.21.17.2 (Assign address)
>> Default gateway 2# ena          (Enable secondary default gateway)
```

9. Enable, apply, and verify the configuration.

```
>> Default gateway 2# ../fwr      (Select the IP Forwarding Menu)
>> IP Forwarding# on             (Turn IP forwarding on)
>> IP Forwarding# apply          (Make your changes active)
>> IP Forwarding# /cfg/l3/cur    (View current IP settings)
```

10. Examine the resulting information. If any settings are incorrect, make the appropriate changes.

11. Save your new configuration changes.

```
>> IP Forwarding # save          (Save for restore after reboot)
```

Using VLANs to segregate broadcast domains

In the previous example, devices that share a common IP network are all in the same broadcast domain. If you want to limit the broadcasts on your network, you could use VLANs to create distinct broadcast domains. For example, as shown in the following procedure, you could create one VLAN for the client trunks, one for the routers, and one for the servers.

In this example, you are adding to the previous configuration.

1. Determine which switch ports and IP interfaces belong to which VLANs.
2. The following table adds port and VLAN information:

Table 22 Subnet routing example: Optional VLAN ports

VLAN	Devices	IP Interface	Switch Port	VLAN #
1	First Floor Client Workstations	2	20	1
	Second Floor Client Workstations	3	21	1
2	Primary Default Router	1	18	2
	Secondary Default Router	1	19	2
3	Common Servers 1	4	1	3
	Common Servers 2	4	2	3

3. Add the switch ports to their respective VLANs.

4. The VLANs shown in the table above are configured as follows:

```
>> # /cfg/l2/vlan 1(Select VLAN 1)
>> VLAN 1# add port 20          (Add port for 1st floor to VLAN 1)
>> VLAN 1# add port 21          (Add port for 2nd floor to VLAN 1)
>> VLAN 1# ena                  (Enable VLAN 1)
>> VLAN 1# ../VLAN 2           (Select VLAN 2)
>> VLAN 2# add port 18          (Add port for default router 1)
>> VLAN 2# add port 19          (Add port for default router 2)
>> VLAN 2# ena                  (Enable VLAN 2)
>> VLAN 2# ../VLAN 3           (Add port for default router 3)
>> VLAN 3# add port 1           (Select VLAN 3)
>> VLAN 3# add port 2           (Select port for common server 1)
>> VLAN 3# ena                  (Enable VLAN 3)
```

5. Each time you add a port to a VLAN, you may get the following prompt:

```
Port 4 is an untagged port and its current PVID is 1.
Confirm changing PVID from 1 to 2 [y/n]?
```

6. Enter y to set the default Port VLAN ID (PVID) for the port.
7. Add each IP interface to the appropriate VLAN.
8. Now that the ports are separated into three VLANs, the IP interface for each subnet must be placed in the appropriate VLAN. The settings are made as follows:

```
>> VLAN 3# /cfg/l3/if 1         (Select IP interface 1 for def. routers)
>> IP Interface 1# vlan 2       (Set to VLAN 2)
>> IP Interface 1# ../if 2      (Select IP interface 2 for first floor)
>> IP Interface 2# vlan 1       (Set to VLAN 1)
>> IP Interface 2# ../if 3      (Select IP interface 3 for second floor)
>> IP Interface 3# vlan 1       (Set to VLAN 1)
>> IP Interface 3# ../if 4      (Select IP interface 4 for servers)
>> IP Interface 4# vlan 3       (Set to VLAN 3)
```

9. Apply and verify the configuration.

```
>> IP Interface 4# apply         (Make your changes active)
>> IP Interface 4# /info/l2/vlan (View current VLAN information)
>> Information# port            (View current port information)
```

10. Examine the resulting information. If any settings are incorrect, make the appropriate changes.
11. Save your new configuration changes.

```
>> Information# save            (Save for restore after reboot)
```

Dynamic Host Configuration Protocol

Dynamic Host Configuration Protocol (DHCP) is a transport protocol that provides a framework for automatically assigning IP addresses and configuration information to other IP hosts or clients in a large TCP/IP network. Without DHCP, the IP address must be entered manually for each network device. DHCP allows a network administrator to distribute IP addresses from a central point and automatically send a new IP address when a device is connected to a different place in the network.

DHCP is an extension of another network IP management protocol, Bootstrap Protocol (BOOTP), with an additional capability of being able to dynamically allocate reusable network addresses and configuration parameters for client operation.

Built on the client/server model, DHCP allows hosts or clients on an IP network to obtain their configurations from a DHCP server, thereby reducing network administration. The most significant configuration the client receives from the server is its required IP address; (other optional parameters include the generic file name to be booted, the address of the default gateway, and so forth).

The DHCP relay agent eliminates the need to have DHCP/BOOTP servers on every subnet. It allows the administrator to reduce the number of DHCP servers deployed on the network and to centralize them. Without the DHCP relay agent, there must be at least one DHCP server deployed at each subnet that has hosts needing to perform the DHCP request.

DHCP relay agent

DHCP is described in RFC 2131, and the DHCP relay agent supported on HP 10GbE switches is described in RFC 1542. DHCP uses UDP as its transport protocol. The client sends messages to the server on port 67 and the server sends messages to the client on port 68.

DHCP defines the methods through which clients can be assigned an IP address for a finite lease period and allowing reassignment of the IP address to another client later. Additionally, DHCP provides the mechanism for a client to gather other IP configuration parameters it needs to operate in the TCP/IP network.

In the DHCP environment, the switch acts as a relay agent. The DHCP relay feature (`/cfg/13/bootp`) enables the switch to forward a client request for an IP address to two BOOTP servers with IP addresses that have been configured on the switch.


When a switch receives a UDP broadcast on port 67 from a DHCP client requesting an IP address, the switch acts as a proxy for the client, replacing the client source IP (SIP) and destination IP (DIP) addresses. The request is then forwarded as a UDP Unicast MAC layer message to two BOOTP servers whose IP addresses are configured on the switch. The servers respond as a UDP Unicast message back to the switch, with the default gateway and IP address for the client. The destination IP address in the server response represents the interface address on the switch that received the client request. This interface address tells the switch on which VLAN to send the server response to the client.

DHCP relay agent configuration

To enable the switch to be the BOOTP forwarder, you need to configure the DHCP/BOOTP server IP addresses on the switch. Generally, you should configure the command on the switch IP interface closest to the client so that the DHCP server knows from which IP subnet the newly allocated IP address should come.

The following figure shows a basic DHCP network example:

Figure 16 DHCP relay agent configuration



In HP 10GbE switch implementation, there is no need for primary or secondary servers. The client request is forwarded to the BOOTP servers configured on the switch. The use of two servers provides failover redundancy. However, no health checking is supported.

Use the following commands to configure the switch as a DHCP relay agent:

```
>> # /cfg/13/bootp
>> Bootstrap Protocol Relay# addr      (Set IP address of BOOTP server)
>> Bootstrap Protocol Relay# addr2    (Set IP address of 2nd BOOTP server)
>> Bootstrap Protocol Relay# on       (Globally turn BOOTP relay on)
>> Bootstrap Protocol Relay# off      (Globally turn BOOTP relay off)
>> Bootstrap Protocol Relay# cur      (Display current configuration)
```

Additionally, DHCP Relay functionality can be assigned on a per interface basis. Use the following command to enable the Relay functionality:

```
>> # /cfg/13/if <1-249>/relay ena
```

Routing Information Protocol

In a routed environment, routers communicate with one another to keep track of available routes. Routers can learn about available routes dynamically, using the Routing Information Protocol (RIP). HP 10GbE switch software supports RIP version 1 (RIPv1) and RIP version 2 (RIPv2) for exchanging TCP/IP route information with other routers.

Distance vector protocol

RIP is known as a distance vector protocol. The vector is the network number and next hop, and the distance is the cost associated with the network number. RIP identifies network reachability based on cost, and cost is defined as hop count. One hop is considered to be the distance from one switch to the next which is typically 1. This cost or hop count is known as the metric.

When a switch receives a routing update that contains a new or changed destination network entry, the switch adds 1 to the metric value indicated in the update and enters the network in the routing table. The IP address of the sender is used as the next hop.

Stability

RIP includes a number of other stability features that are common to many routing protocols. For example, RIP implements the split horizon and hold-down mechanisms to prevent incorrect routing information from being propagated.

RIP prevents routing loops from continuing indefinitely by implementing a limit on the number of hops allowed in a path from the source to a destination. The maximum number of hops in a path is 15. The network destination network is considered unreachable if increasing the metric value by 1 causes the metric to be 16 (that is infinity). This limits the maximum diameter of a RIP network to less than 16 hops.

RIP is often used in stub networks and in small autonomous systems that do not have many redundant paths.

Routing updates

RIP sends routing-update messages at regular intervals and when the network topology changes. Each router advertises routing information by sending a routing information update every 30 seconds. If a router doesn't receive an update from another router for 180 seconds, those routes provided by that router are declared invalid. After another 120 seconds without receiving an update for those routes, the routes are removed from the routing table and respective regular updates.

When a router receives a routing update that includes changes to an entry, it updates its routing table to reflect the new route. The metric value for the path is increased by 1, and the sender is indicated as the next hop. RIP routers maintain only the best route (the route with the lowest metric value) to a destination.

For more information see The Configuration Menu, Routing Information Protocol Configuration (/cfg/13/rip) in the *HP 10Gb Ethernet BL-c Switch Command Reference*.

RIPv1

RIP version 1 use broadcast User Datagram Protocol (UDP) data packets for the regular routing updates. The main disadvantage is that the routing updates do not carry subnet mask information. Hence, the router cannot determine whether the route is a subnet route or a host route. It is of limited usage after the introduction of RIPv2. For more information about RIPv1 and RIPv2, refer to RFC 1058 and RFC 2453.

RIPv2

RIPv2 is the most popular and preferred configuration for most networks. RIPv2 expands the amount of useful information carried in RIP messages and provides a measure of security. For a detailed explanation of RIPv2, refer to RFC 1723 and RFC 2453.

RIPv2 improves efficiency by using multicast UDP (address 224.0.0.9) data packets for regular routing updates. Subnet mask information is provided in the routing updates. A security option is added for authenticating routing updates, by using a shared password. HP 10GbE switch software supports using clear password for RIPv2.

RIPv2 in RIPv1 compatibility mode

HP 10GbE switch software allows you to configure RIPv2 in RIPv1 compatibility mode, for using both RIPv2 and RIPv1 routers within a network. In this mode, the regular routing updates use broadcast UDP data packet to allow RIPv1 routers to receive those packets. With RIPv1 routers as recipients, the routing updates have to carry natural or host mask. Hence, it is not a recommended configuration for most network topologies.

NOTE: When using both RIPv1 and RIPv2 within a network, use a single subnet mask throughout the network.

RIP Features

HP 10GbE switch software provides the following features to support RIPv1 and RIPv2:

Poison

Simple split horizon in RIP scheme omits routes learned from one neighbor in updates sent to that neighbor. That is the most common configuration used in RIP that is setting this Poison to `disable`. Split horizon with poisoned reverse includes such routes in updates, but sets their metrics to 16. The disadvantage of using this feature is the increase of size in the routing updates.

Triggered updates

Triggered updates are an attempt to speed up convergence. When Triggered Updates is enabled (`/cfg/l3/rip/if x/trigg ena`), whenever a router changes the metric for a route, it sends update messages almost immediately, without waiting for the regular update interval. It is recommended to enable Triggered Updates.

Multicast

RIPv2 messages use IP multicast address (224.0.0.9) for periodic broadcasts. Multicast RIPv2 announcements are not processed by RIPv1 routers. IGMP is not needed since these are inter-router messages which are not forwarded.

To configure RIPv2 in RIPv1-compatibility mode, set `multicast` to `disable`.

Default

The RIP router can listen and supply a default route, usually represented as 0.0.0.0 in the routing table. When a router does not have an explicit route to a destination network in its routing table, it uses the default route to forward those packets.

Metric

The metric field contains a configurable value between 1 and 15 (inclusive) which specifies the current metric for the interface. The metric value typically indicates the total number of hops to the destination. The metric value of 16 represents an unreachable destination.

Authentication

RIPv2 authentication uses plaintext password for authentication. If configured using `Authentication password`, then it is necessary to enter an authentication key value.

The following method is used to authenticate a RIP message:

- If the router is not configured to authenticate RIPv2 messages, then RIPv1 and unauthenticated RIPv2 messages are accepted; authenticated RIPv2 messages are discarded.
- If the router is configured to authenticate RIPv2 messages, then RIPv1 messages and RIPv2 messages which pass authentication testing are accepted; unauthenticated and failed authentication RIPv2 messages are discarded.

For maximum security, RIPv1 messages are ignored when authentication is enabled (`cfg/l3/rip/if x/auth password`); otherwise, the routing information from authenticated messages is propagated by RIPv1 routers in an unauthenticated manner.

RIP configuration example

NOTE: An interface RIP disabled uses all the default values of the RIP, no matter how the RIP parameters are configured for that interface. RIP sends out RIP regular updates to include an Up interface, but not a Down interface.

1. Add VLANs for routing interfaces.

```
>> Main# cfg/l2/vlan 2/ena (Enable VLAN 2)
>> VLAN 2# add 20 (Add port 20 to VLAN 2)

Port 20 is an UNTAGGED port and its current PVID is 1.
Confirm changing PVID from 1 to 2 [y/n]: y

>> VLAN 2# /cfg/l2/vlan 3/ena (Enable VLAN 3)
>> VLAN 3# add 21 (Add port 21 to VLAN 3)

Port 21 is an UNTAGGED port and its current PVID is 1.
Confirm changing PVID from 1 to 3 [y/n]: y
```

2. Add IP interfaces to VLANs.

```
>> Main# cfg/l3/if 2/ena (Enable interface 2)
>> IP Interface 2# addr 102.1.1.1 (Define IP address for interface 2)
>> IP Interface 2# vlan 2 (Add interface 2 to VLAN 2)
>> IP Interface 2# /cfg/l3/if 3/ena (Enable interface 3)
>> IP Interface 3# addr 103.1.1.1 (Define IP address for interface 3)
>> IP Interface 3# vlan 3 (Add interface 3 to VLAN 3)
```

3. Turn on RIP globally and enable RIP for each interface. IP Forwarding must be on (/cfg/l3/frwd/on) before you turn RIP on.

```
>> Main# cfg/l3/rip on (Turn on RIP globally)
>> Routing Information Protocol# if 2/ena (Enable RIP on IP interface 2)
>> RIP Interface 2# ..
>> Routing Information Protocol# if 3/ena (Enable RIP on IP interface 3)
>> RIP Interface 3# apply (Apply your changes)
>> RIP Interface 3# save (Save the configuration)
```

Use the /maint/route/dump command to check the current valid routes in the routing table of the switch.

For those RIP learned routes, within the garbage collection period, that are routes phasing out of the routing table with metric 16, use the /info/l3/rip/routes command. Locally configured static routes do not appear in the RIP Routes table.

IGMP Snooping

Introduction

IGMP Snooping allows the switch to forward multicast traffic only to those ports that request it. IGMP Snooping prevents multicast traffic from being flooded to all data ports. The switch learns which server hosts are interested in receiving multicast traffic, and forwards it only to ports connected to those servers.

The following topics are discussed in this chapter:

- Overview
- IGMPv3
- FastLeave
- IGMP Filtering
- Static Multicast Router
- IGMP Snooping Configuration example

Overview

Internet Group Management Protocol (IGMP) is used by IP Multicast routers to learn about the existence of host group members on their directly attached subnet (see RFC 2236). The IP Multicast routers get this information by broadcasting IGMP Query Reports and listening for IP hosts reporting their host group memberships. This process is used to set up a client/server relationship between an IP Multicast source that provides the data streams and the clients that want to receive the data.

IGMP Snooping conserves bandwidth. With IGMP Snooping, the switch learns which ports are interested in receiving multicast data, and forwards multicast data only to those ports. In this way, other ports are not burdened with unwanted multicast traffic.

The switch currently supports snooping for IGMP version 1, version 2, and version 3.

The switch can sense IGMP Membership Reports from attached host servers and act as a proxy to set up a dedicated path between the requesting host and a local IP Multicast router. After the pathway is established, the switch blocks the IP Multicast stream from flowing through any port that does not connect to a host member, thus conserving bandwidth.

The client-server path is set up as follows:

- An IP Multicast Router (Mrouter) sends Membership Queries to the switch, which forwards them to all ports in a given VLAN.
- Hosts that want to receive the multicast data stream send Membership Reports to the switch, which sends a proxy Membership Report to the Mrouter.
- The switch sets up a path between the Mrouter and the host, and blocks all other ports from receiving the multicast.
- Periodically, the Mrouter sends Membership Queries to ensure that the host wants to continue receiving the multicast. If the host fails to respond with a Membership Report, the Mrouter stops sending the multicast to that path.

- The host can send an IGMPv2 Leave report to the switch, which sends a proxy Leave report to the Mrouter. The multicast path is terminated immediately.

A maximum of 8 VLANs can be configured for IGMP Snooping. The switch can learn up to 16 multicast routers, and supports up to 1,000 multicast groups.

IGMPv3

IGMPv3 includes new membership report messages to extend IGMP functionality. The switch provides snooping capability for all types of IGMP version 3 (IGMPv3) Membership Reports, as described in RFC 3376.

IGMPv3 supports Source-Specific Multicast (SSM). SSM identifies session traffic by both source and group addresses. The switch uses *source filtering*, which allows hosts to report interest in receiving multicast packets only from specific source addresses, or from all but specific source addresses.

The switch supports the following IGMPv3 filter modes:

- INCLUDE mode: The host requests membership to a multicast group and provides a list of IP addresses from which it wants to receive traffic.
- EXCLUDE mode: The host requests membership to a multicast group and provides a list of IP addresses from which it *does not* want to receive traffic. This indicates that the host wants to receive traffic only from sources that are not part of the Exclude list.

To disable snooping on EXCLUDE mode reports, use the following command:

```
/cfg/13/igmp/snoop/igmpv3/exclude dis
```

By default, the switch snoops the first eight sources listed in the IGMPv3 Group Record. Use the following command to change the number of snooping sources:

```
/cfg/13/igmp/snoop/igmpv3/sources <1-64>
```

IGMPv3 Snooping is compatible with IGMPv1 and IGMPv2 Snooping. You can disable snooping on version 1 and version 2 reports, using the following command:

```
/cfg/13/igmp/snoop/igmpv3/v1v2 dis
```

FastLeave

When the switch with IGMP Snooping enabled receives an IGMPv2 leave message, it sends a Group-Specific Query to determine if any other devices in the same group (and on the same port) are still interested in the specified multicast group traffic. The switch removes the affiliated port from that particular group, if the following conditions apply:

- If the switch does not receive an IGMP Membership Report message within the query-response-interval
- If no multicast routers have been learned on that port.

With FastLeave enabled on the VLAN, a port can be removed immediately from the port list of the group entry when the IGMP Leave message is received, unless a multicast router was learned on the port.

Enable FastLeave only on VLANs that have only one host connected to each physical port.

IGMP Filtering

With IGMP Filtering, you can allow or deny a port to send and receive multicast traffic to certain multicast groups. Unauthorized users are restricted from streaming multicast traffic across the network.

If access to a multicast group is denied, IGMP Membership Reports from the port for that group are dropped, and the port is not allowed to receive IP multicast traffic from that group. If access to the multicast group is allowed, Membership Reports from the port are forwarded for normal processing.

To configure IGMP Filtering, you must globally enable IGMP Filtering, define an IGMP Filter, assign the filter to a port, and enable IGMP Filtering on the port. To define an IGMP Filter, you must configure a range of IP multicast groups, choose whether the filter will allow or deny multicast traffic for groups within the range, and enable the filter.

NOTE: Low-numbered filters take precedence over high-number filters. For example, the action defined for IGMP Filter 1 supersedes the action defined for IGMP Filter 2.

Configuring the range

Each IGMP Filter allows you to set a start and end point that defines the range of IP addresses upon which the filter takes action. Each IP address in the range must be between 224.0.0.0 and 239.255.255.255.

Configuring the action

Each IGMP Filter can allow or deny IP multicasts to the range of IP addresses configured. If you configure the filter to deny IP multicasts, then IGMP Membership Reports from multicast groups within the range are dropped.

You can configure a secondary filter to allow IP multicasts to a small range of addresses within a larger range that a primary filter is configured to deny. The two filters work together to allow IP multicasts to a small subset of addresses within the larger range of addresses. The secondary filter must have a lower number than the primary filter, so that it takes precedence.

Static multicast router

A static multicast router (Mrouter) can be configured for a particular port on a particular VLAN. A static Mrouter does not have to be learned through IGMP Snooping.

You can configure static Mrouters on any switch port except the management port 17. The switch supports up to total of sixteen static Mrouters.

When you configure a static Mrouter on a VLAN, it replaces any dynamic Mrouters learned through IGMP Snooping.

IGMP Snooping configuration example

This section provides steps to configure IGMP Snooping on the switch, using the Command Line Interface (CLI) or the Browser-based Interface (BBI).

Configuring IGMP Snooping (CLI example)

1. Configure port and VLAN membership on the switch, as described in the “Configuring ports and VLANs (CLI example)” section in the “VLANs” chapter.
2. Add VLANs to IGMP Snooping and enable the feature.

```
>> /cfg/l3/igmp/on                (Globally turn IGMP on)
>> IGMP# snoop                    (Select IGMP Snooping menu)
>> IGMP Snoop# add 1              (Add VLAN 1 to IGMP Snooping)
>> IGMP Snoop# apply              (Make your changes active)
```

3. Enable IGMPv3 Snooping (optional).

```
>> IGMP Snoop# igmpv3            (Select IGMPv3 menu)
>> IGMP V3 Snoop# ena           (Enable IGMPv3 Snooping)
```

4. Apply and save the configuration.

```
>> IGMP V3 Snoop# apply          (Apply the configuration)
>> IGMP V3 Snoop# save           (Save your changes)
```

5. View dynamic IGMP information.

```
>> /info/l3/igmp                (Select IGMP Information menu)
>> IGMP Multicast# dump         (Show IGMP Group information)

Note: Local groups (224.0.0.x) are not snooped and will not appear.
Source          Group          VLAN  Port   Version  Mode Expires  Fwd
-----
224.10.2.0      232.1.0.0      2     18     V3       INC  4:17     Yes
228.12.1.0      232.1.0.4      2     18     V3       INC  4:18     Yes
*               235.2.0.1      3     21     V3       INC  -         No
*               236.1.0.2      3     21     V3       EXC  -         Yes

>> /info/l3/igmp/mrouter        (Access Mrouter information menu)
>> Mrouter# dump                (Show IGMP Group information)

VLAN  Port   Version Expires  Max Query Resp. Time QRV QVIC
-----
1     21     V2     static  unknown  -     -
2     20     V3     4:09   128      2     125
```

These commands display information about IGMP Groups and Mrouters learned through IGMP Snooping.

Configuring IGMP Filtering (CLI example)

1. Enable IGMP Filtering on the switch.

```
>> /cfg/l3/igmp/igmpflt          (Select IGMP Filtering menu)
>> IGMP Filter# ena              (Enable IGMP Filtering)
Current status: disabled
New status: enabled
```

2. Define an IGMP Filter.

```
>> //cfg/l3/igmp/igmpflt        (Select IGMP Filtering menu)
>>IGMP Filter# filter 1         (Select Filter 1 Definition menu)
>>IGMP Filter 1 Definition# range 224.0.1.0 (Enter first IP
                                     address of the range)

Current multicast address2:
Enter new multicast address2: 226.0.0.0 (Enter second IP
                                     address of the range)

Current multicast address1:
New pending multicast address1: 224.0.1.0
Current multicast address2:
New pending multicast address2: 226.0.0.0
>>IGMP Filter 1 Definition# action deny (Deny multicast traffic)
>>IGMP Filter 1 Definition# ena      (Enable the filter)
```

3. Assign the IGMP Filter to a port.

```
>> //cfg/l3/igmp/igmpflt        (Select IGMP Filtering menu)
>>IGMP Filter# port 21          (Select port 21)
>>IGMP Port 21# filt ena        (Enable IGMP Filtering on the port)
Current port 21 filtering: disabled
New port 21 filtering: enabled
>>IGMP Port 21# add 1           (Add IGMP Filter 1 to the port)
>>IGMP Port 21# apply           (Make your changes active)
```

Configuring a Static Mrouter (CLI example)

1. Configure a port to which the static Mrouter is connected, and enter the appropriate VLAN.

```
>> /cfg/l3/igmp/mrouter         (Select IGMP Mrouter menu)
>> Static Multicast Router# add 20 (Add port 20 as Static
                                     Mrouter port)


Enter VLAN number: (1-4094) 1      (Enter the VLAN number)
Enter the version number of mrouter [1|2|3]: 2 (Enter the IGMP
                                               version number)
```

2. Apply, verify, and save the configuration.

```
>> Static Multicast Router# apply (Apply the configuration)
>> Static Multicast Router# cur   (View the configuration)
>> Static Multicast Router# save  (Save the configuration)
```

Configuring IGMP Snooping (BBI example)

1. Configure port and VLAN membership on the switch, as described in the “Configuring ports and VLANs (BBI example)” section in the “VLANs” chapter.
2. Configure IGMP Snooping.
 - a. Click the **Configure** context button.
 - b. Open the IGMP folder, and select **IGMP Snooping** (click the underlined text, not the folder).



- c. Enable IGMP Snooping.

IGMP Snooping Configuration

IGMP on ?	on ▾
Set report timeout	10
Set multicast router timeout	255
Set robust value or expected packet loss on subnet	2
Set query interval	125
Aggregate IGMP report	enabled ▾
Flood unregistered IPMC	disabled ▾
Set Source IP for GSQ proxy	255.255.255.255
Remove all VLAN(s) from IGMP Snooping	none ▾
IGMP V3 snooping on ?	Disabled ▾
Set number of sources to snoop in GR	8
Exclude ?	Enabled ▾
v1v2 ?	Enabled ▾

Configured VLANs

VLAN ID:#

Add>>

<<Remove

Snooping VLANs

VLAN ID:#
VLAN:1

VLANs without Fastleave

VLAN ID:#
VLAN:1

Add>>

<<Remove


VLANs with Fastleave

VLAN ID:#

Submit

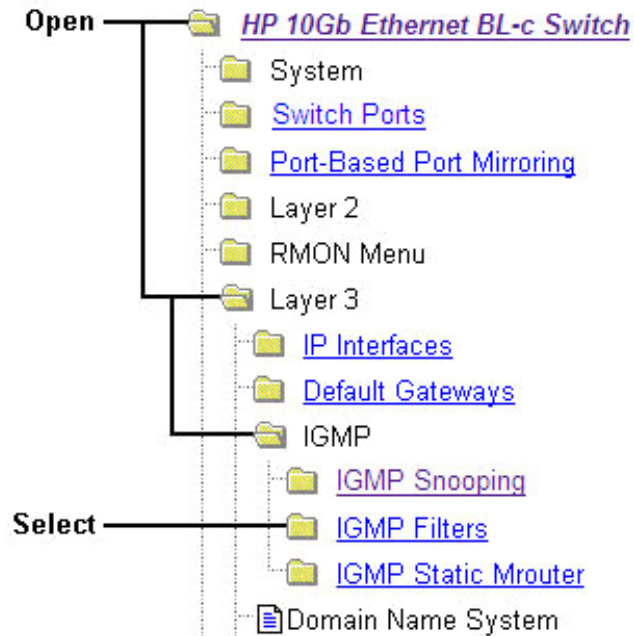
Submit Reset

- d. Click **Submit**.
- 3. Apply, verify, and save the configuration.



Configuring IGMP Filtering (BBI example)

1. Configure IGMP Snooping.
2. Enable IGMP Filtering.
 - a. Click the **Configure** context button.
 - b. Open the IGMP folder, and select **IGMP Filters** (click the underlined text, not the folder).



- c. Enable IGMP Filtering globally.

IGMP Filters Configuration

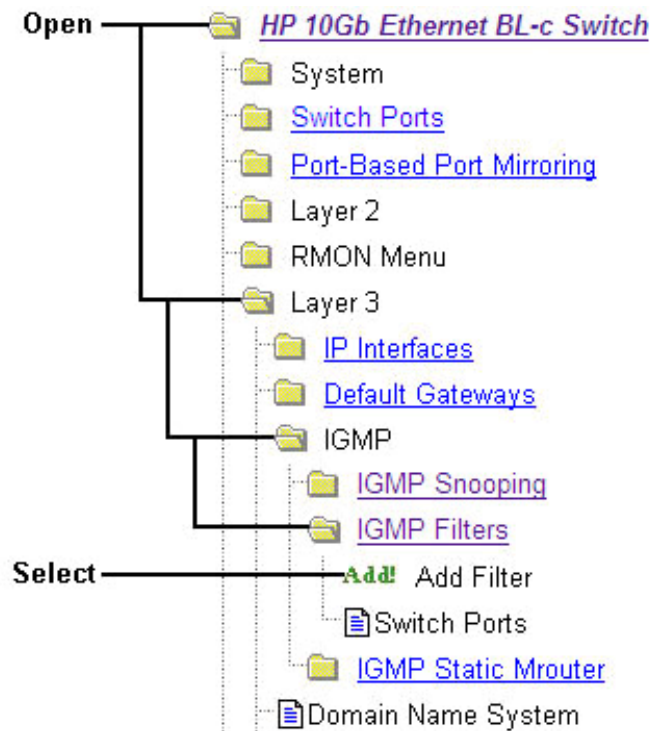
IGMP Filter Enabled?

Filter ID	Enabled?	Action	Range
<u>1</u>	ena	deny	224.0.1.0- 226.0.0.0

- d. Click **Submit**.

3. Define the IGMP Filter.

- a. Select **Layer 3 > IGMP > IGMP Filters > Add Filter**.



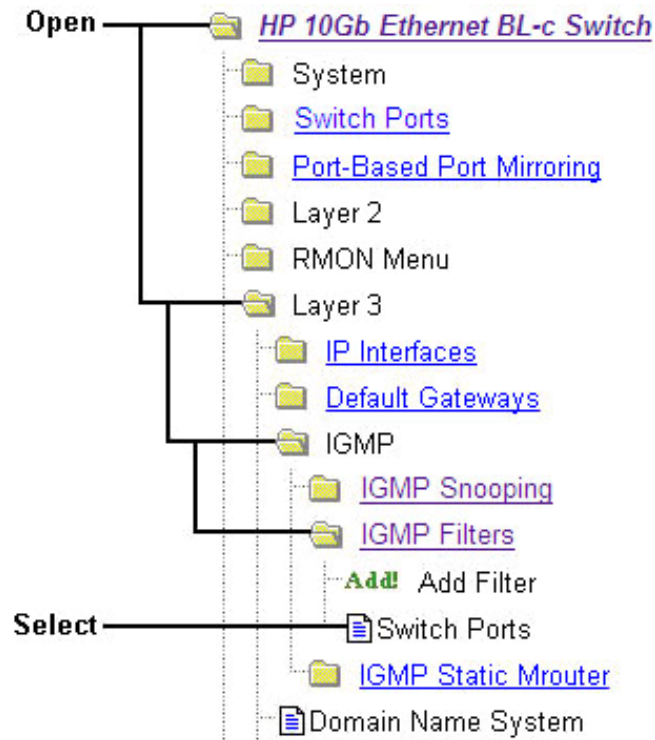
- b. Enable the IGMP Filter. Assign the range of IP multicast addresses and the filter action (allow or deny).

IGMP Filter Configuration

Filter Identifier (1 - 16)	<input type="text" value="1"/>
Enabled?	<input type="button" value="Enabled"/>
Range 1 IP Multicast Address	<input type="text" value="224.0.1.0"/>
Range 2 IP Multicast Address	<input type="text" value="226.0.0.0"/>
Action	<input type="button" value="Deny"/>

- c. Click **Submit**.

4. Assign the filter to a port and enable IGMP Filtering on the port.
 - a. Select **Layer 3 > IGMP > IGMP Filters > Switch Ports**.



- b. Select a port from the list.

IGMP Filtering Port Configuration	
Switch Port	IGMP Filter Processing?
1	disabled
2	disabled
3	disabled
4	disabled
⋮	
18	disabled
19	disabled
20	disabled
21	disabled

Select ———→

- c. Enable IGMP Filtering on the port. Select a filter in the IGMP Filters Available list, and click **Add**.

IGMP Filtering - Port 21 Configuration

Enable/Disable Filtering on Port: enabled

IGMP Filters Available

Filter ID

IGMP Filters Selected


Filter ID
1

Add >>

<< Remove

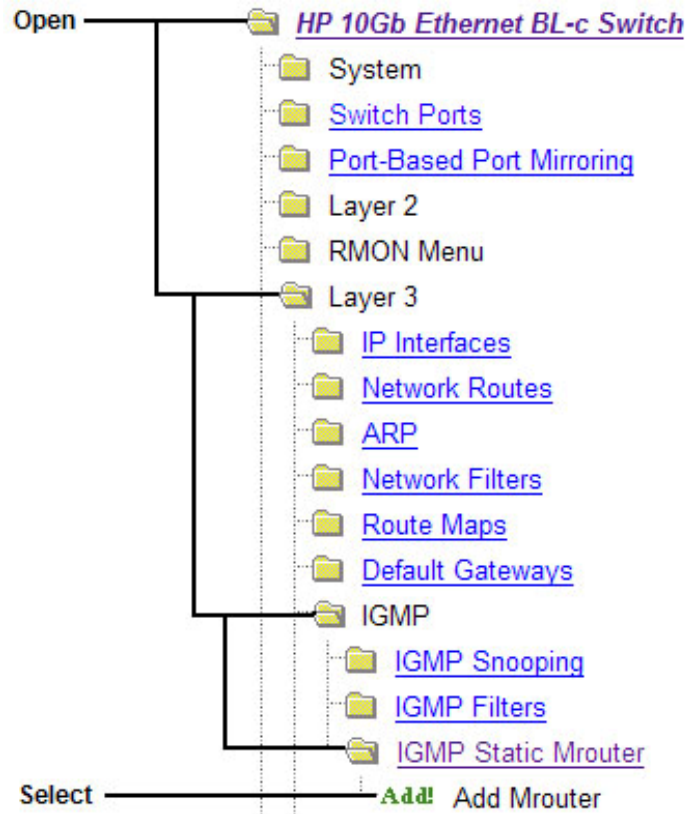
Submit

- d. Click **Submit**.
5. Apply, verify, and save the configuration.



Configuring a Static Multicast Router (BBI example)

1. Configure Static Mrouter.
 - a. Click the **Configure** context button.
 - b. Open the Switch folder and select **Layer 3 > IGMP > IGMP Static Mrouter > Add Mrouter**.



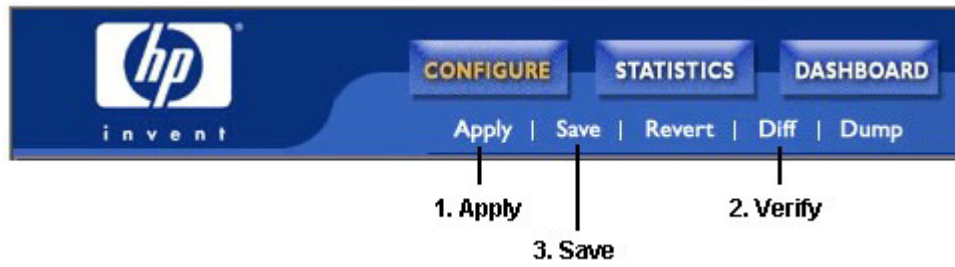
- c. Enter a port number, VLAN ID number, and IGMP version number.

Static Multicast Router Configuration for Port

Mrouter Port ID (egs 1, 14..)	<input type="text" value="20"/>
Vlan ID	<input type="text" value="1"/>
IGMP Version?	Version2 ▾

- d. Click **Submit**.

2. Apply, verify, and save the configuration.



OSPF

The HP 10GbE switch software supports the Open Shortest Path First (OSPF) routing protocol. The switch implementation conforms to the OSPF version 2 specifications detailed in Internet RFC 1583. The following sections discuss OSPF support for the HP 10GbE switch:

- **OSPF Overview:** This section provides information on OSPF concepts, such as types of OSPF areas, types of routing devices, neighbors, adjacencies, link state database, authentication, and internal versus external routing.
- **OSPF Implementation in HP 10GbE switch software.** This section describes how OSPF is implemented in switch software, such as configuration parameters, electing the designated router, summarizing routes, defining route maps and so forth.
- **OSPF Configuration Examples.** This section provides step-by-step instructions on configuring different configuration examples:
 - Creating a simple OSPF domain
 - Creating virtual links
 - Summarizing routes

OSPF overview

OSPF is designed for routing traffic within a single IP domain called an Autonomous System (AS). The AS can be divided into smaller logical units known as areas.

All routing devices maintain link information in their own Link State Database (LSDB). The LSDB for all routing devices within an area is identical but is not exchanged between different areas. Only routing updates are exchanged between areas, thereby significantly reducing the overhead for maintaining routing information on a large, dynamic network.

The following sections describe key OSPF concepts.


Types of OSPF areas

An AS can be broken into logical units known as areas. In any AS with multiple areas, one area must be designated as area 0, known as the backbone. The backbone acts as the central OSPF area. All other areas in the AS must be connected to the backbone. Areas inject summary routing information into the backbone, which then distributes it to other areas as needed.

OSPF defines the following types of areas:

- **Stub Area**—an area that is connected to only one other area. External route information is not distributed into stub areas.
- **Not-So-Stubby-Area (NSSA)**—similar to a stub area with additional capabilities. External routes from outside the AS can be advertised within the NSSA but only if they are originated by a router from inside the NSSA area. External routes originated by a router that is not a member of the NSSA area are not advertised in the NSSA area.
- **Transit Area**—an area that allows area summary information to be exchanged between routing devices. The backbone (area 0), any area that contains a virtual link to connect two areas, and any area that is not a stub area or an NSSA are considered transit areas

Figure 17 OSPF area types

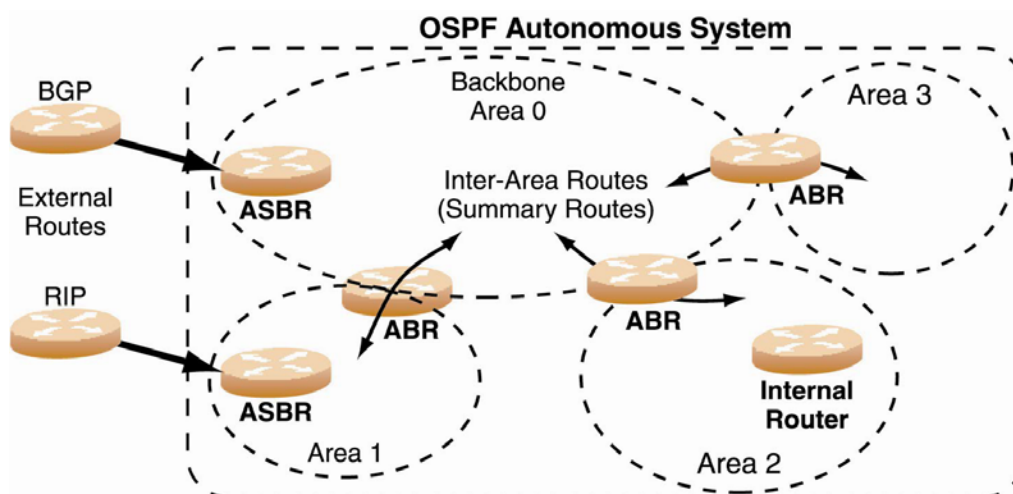


Types of OSPF routing devices

As shown in the figure, OSPF uses the following types of routing devices:

- Internal Router (IR)—a router that has all of its interfaces within the same area. IRs maintain LSDBs identical to those of other routing devices within the local area.
- Area Border Router (ABR)—a router that has interfaces in multiple areas. ABRs maintain one LSDB for each connected area and disseminate routing information between areas.
- Autonomous System Boundary Router (ASBR)—a router that acts as a gateway between the OSPF domain and non-OSPF domains, such as RIP, BGP, and static routes.

Figure 18 OSPF domain and an autonomous system



Neighbors and adjacencies

In areas with two or more routing devices, neighbors and adjacencies are formed.

Neighbors are routing devices that maintain information about each others' health. To establish neighbor relationships, routing devices periodically send hello packets on each of their interfaces. All routing devices that share a common network segment, appear in the same area, and have the same health parameters (hello and dead intervals) and authentication parameters respond to each other's hello packets and become neighbors. Neighbors continue to send periodic hello packets to advertise their health to neighbors. In turn, they listen to hello packets to determine the health of their neighbors and to establish contact with new neighbors.

The hello process is used for electing one of the neighbors as the area's Designated Router (DR) and one as the area's Backup Designated Router (BDR). The DR is adjacent to all other neighbors and acts as the central contact for database exchanges. Each neighbor sends its database information to the DR, which relays the information to the other neighbors.

The BDR is adjacent to all other neighbors (including the DR). Each neighbor sends its database information to the BDR just as with the DR, but the BDR merely stores this data and does not distribute it. If the DR fails, the BDR will take over the task of distributing database information to the other neighbors.

Link-State Database

OSPF is a link-state routing protocol. A link represents an interface (or routable path) from the routing device. By establishing an adjacency with the DR, each routing device in an OSPF area maintains an identical Link-State Database (LSDB) describing the network topology for its area.

Each routing device transmits a Link-State Advertisement (LSA) on each of its interfaces. LSAs are entered into the LSDB of each routing device. OSPF uses flooding to distribute LSAs between routing devices.

When LSAs result in changes to the routing device's LSDB, the routing device forwards the changes to the adjacent neighbors (the DR and BDR) for distribution to the other neighbors.

OSPF routing updates occur only when changes occur, instead of periodically. For each new route, if an adjacency is interested in that route (for example, if configured to receive static routes and the new route is indeed static), an update message containing the new route is sent to the adjacency. For each route removed from the route table, if the route has already been sent to an adjacency, an update message containing the route to withdraw is sent.

Shortest Path First Tree

The routing devices use a link-state algorithm (Dijkstra's algorithm) to calculate the shortest path to all known destinations, based on the cumulative cost required to reach the destination.

The cost of an individual interface in OSPF is an indication of the overhead required to send packets across it. The cost is inversely proportional to the bandwidth of the interface. A lower cost indicates a higher bandwidth.

Internal versus external routing

To ensure effective processing of network traffic, every routing device on your network needs to know how to send a packet (directly or indirectly) to any other location/destination in your network. This is referred to as internal routing and can be done with static routes or using active internal routing protocols, such as OSPF, RIP, or RIPv2.

It is also useful to tell routers outside your network (upstream providers or peers) about the routes you have access to in your network. Sharing of routing information between autonomous systems is known as external routing.

Typically, an AS will have one or more border routers (peer routers that exchange routes with other OSPF networks) as well as an internal routing system enabling every router in that AS to reach every other router and destination within that AS.

When a routing device advertises routes to boundary routers on other autonomous systems, it is effectively committing to carry data to the IP space represented in the route being advertised. For example, if the routing device advertises 192.204.4.0/24, it is declaring that if another router sends data destined for any address in the 192.204.4.0/24 range, it will carry that data to its destination.

OSPF implementation in HP 10GbE switch software

The HP 10GbE switch supports a single instance of OSPF and up to 4 K routes on the network. The following sections describe OSPF implementation in switch software:

- Configurable Parameters
- Defining Areas
- Interface Cost
- Electing the Designated Router and Backup
- Summarizing Routes
- Default Routes
- Virtual Links
- Router ID
- Authentication

Configurable parameters

In HP 10GbE switch software, OSPF parameters can be configured through the Command Line Interface (CLI), Browser-Based Interface (BBI) for HP 10GbE switches, or through SNMP. For more information, see “Accessing the Switch.”

The CLI supports the following parameters: interface output cost, interface priority, dead and hello intervals, retransmission interval, and interface transit delay.

OSPF traps—Traps produce messages upon certain events or error conditions, such as missing a hello, failing a neighbor, or recalculating the SPF.

In addition to the above parameters, you can also specify the following:

- Link-State Database size—The size of the external LSA database can be specified to help manage the memory resources on the switch.
- Shortest Path First (SPF) interval—Time interval between successive calculations of the shortest path tree using the Dijkstra’s algorithm.

- Stub area metric—A stub area can be configured to send a numeric metric value such that all routes received via that stub area carry the configured metric to potentially influence routing decisions.
- Default routes—Default routes with weight metrics can be manually injected into transit areas. This helps establish a preferred route when multiple routing devices exist between two areas. It also helps route traffic to external networks.

Defining areas

If you are configuring multiple areas in your OSPF domain, one of the areas must be designated as area 0, known as the backbone. The backbone is the central OSPF area and is usually physically connected to all other areas. The areas inject routing information into the backbone which, in turn, disseminates the information into other areas.

Since the backbone connects the areas in your network, it must be a contiguous area. If the backbone is partitioned (possibly as a result of joining separate OSPF networks), parts of the AS will be unreachable, and you will need to configure virtual links to reconnect the partitioned areas (see “Virtual Links”).

Up to three OSPF areas can be connected to the HP 10GbE switch. To configure an area, the OSPF number must be defined and then attached to a network interface on the switch. The full process is explained in the following sections.

An OSPF area is defined by assigning two pieces of information—an area index and an area ID. The command to define an OSPF area is as follows:

```
>> # /cfg/l3/ospf/aindex <area index>/areaid <n.n.n.n>
```

NOTE: The `aindex` option above is an arbitrary index used only on the switch and does not represent the actual OSPF area number. The actual OSPF area number is defined in the `areaid` portion of the command as explained in the following sections.

Assigning the area index

The `aindex <area index>` option is actually just an arbitrary index (0-2) used only by the switch. This index does not necessarily represent the OSPF area number, though for configuration simplicity, it should where possible.

For example, both of the following sets of commands define OSPF area 0 (the backbone) and area 1 because that information is held in the area ID portion of the command. However, the first set of commands is easier to maintain because the arbitrary area indexes agree with the area IDs:

- Area index and area ID agree
 - `/cfg/l3/ospf/aindex 0/areaid 0.0.0.0` (Use index 0 to set area 0 in ID octet format)
 - `/cfg/l3/ospf/aindex 1/areaid 0.0.0.1` (Use index 1 to set area 1 in ID octet format)
- Area index set to an arbitrary value
 - `/cfg/l3/ospf/aindex 1/areaid 0.0.0.0` (Use index 1 to set area 0 in ID octet format)
 - `/cfg/l3/ospf/aindex 2/areaid 0.0.0.1` (Use index 2 to set area 1 in ID octet format)

Using the area ID to assign the OSPF area number

The OSPF area number is defined in the `areaid <IP address>` option. The octet format is used in order to be compatible with two different systems of notation used by other OSPF network vendors. There are two valid ways to designate an area ID:

- Placing the area number in the last octet (0.0.0.*n*)
Most common OSPF vendors express the area ID number as a single number. For example, the Cisco IOS-based router command `network 1.1.1.0 0.0.0.255 area 1` defines the area number simply as `area 1`. On the switch, using the last octet in the area ID, `area 1` is equivalent to `areaid 0.0.0.1`.
- Multi-octet (*IP address*)
Some OSPF vendors express the area ID number in multi-octet format. For example, `area 2.2.2.2` represents OSPF area 2 and can be specified directly on the switch as `areaid 2.2.2.2`.

NOTE: Although both types of area ID formats are supported, be sure that the area IDs are in the same format throughout an area.

Attaching an area to a network

Once an OSPF area has been defined, it must be associated with a network. To attach the area to a network, you must assign the OSPF area index to an IP interface that participates in the area. The format for the command is as follows:

```
>> # /cfg/l3/ospf/if <interface number>/aindex <area index>
```

For example, the following commands could be used to configure IP interface 14 for a presence on the 10.10.10.1/24 network, to define OSPF area 1, and to attach the area to the network:

```
>> # /cfg/l3/if 14          (Select menu for IP interface 14)
>> IP Interface 14# addr 10.10.10.1(Define IP address on backbone
                           network)
>> IP Interface 14# mask 255.255.255.0(Define IP mask on backbone)
>> IP Interface 14# ena          (Enable IP interface 14)
>> IP Interface 14# ../ospf/aindex 1(Select menu for area index 1)
>> OSPF Area (index) 1 # areaid 0.0.0.1(Define area ID as OSPF area 1)
>> OSPF Area (index) 1 # ena      (Enable area index 1)
>> OSPF Area (index) 1 # ../if 14 (Select OSPF menu for interface 14)
>> OSPF Interface 14# aindex 1    (Attach area to network on
                                   interface 14)
>> OSPF Interface 14# enable      (Enable interface 14 for area index 1)
```

Interface cost

The OSPF link-state algorithm (Dijkstra's algorithm) places each routing device at the root of a tree and determines the cumulative cost required to reach each destination. Usually, the cost is inversely proportional to the bandwidth of the interface. Low cost indicates high bandwidth. You can manually enter the cost for the output route with the following command:

```
>> # /cfg/l3/ospf/if <OSPF interface number>/cost <cost value (1-65535)>
```


Electing the designated router and backup

In any area with more than two routing devices, a Designated Router (DR) is elected as the central contact for database exchanges among neighbors, and a Backup Designated Router (BDR) is elected in case the DR fails.

DR and BDR elections are made through the hello process. The election can be influenced by assigning a priority value to the OSPF interfaces on the switch. The command is as follows:

```
>># /cfg/13/ospf/if <OSPF interface number>/prio <priority value (0-255)>
```

A priority value of 255 is the highest, and 1 is the lowest. A priority value of 0 specifies that the interface cannot be used as a DR or BDR. In case of a tie, the routing device with the lowest router ID wins.

Summarizing routes

Route summarization condenses routing information. Without summarization, each routing device in an OSPF network would retain a route to every subnet in the network. With summarization, routing devices can reduce some sets of routes to a single advertisement, reducing both the load on the routing device and the perceived complexity of the network. The importance of route summarization increases with network size.

Summary routes can be defined for up to 16 IP address ranges using the following command:

```
>> # /cfg/13/ospf/range <range number>/addr <IP address>/mask <mask>
```

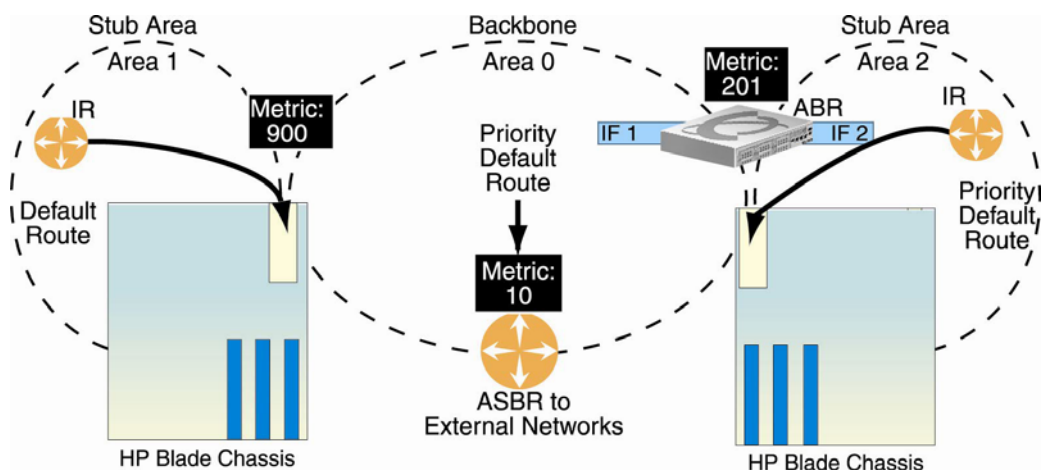
where *<range number>* is a number 1 to 16, *<IP address>* is the base IP address for the range, and *<mask>* is the IP address mask for the range.

Default routes

When an OSPF routing device encounters traffic for a destination address it does not recognize, it forwards that traffic along the default route. Typically, the default route leads upstream toward the backbone until it reaches the intended area or an external router.

Each switch acting as an ABR automatically inserts a default route into each attached area. In simple OSPF stub areas or NSSAs with only one ABR leading upstream (see Area 1 in the figure below), any traffic for IP address destinations outside the area is forwarded to the switch's IP interface, and then into the connected transit area (usually the backbone). Since this is automatic, no further configuration is required for such areas.

Figure 19 Injecting default routes



In more complex OSPF areas with multiple ABRs or ASBRs (such as area 0 and area 2 in the figure), there are multiple routes leading from the area. In such areas, traffic for unrecognized destinations cannot tell which route leads upstream without further configuration.

To resolve the situation and select one default route among multiple choices in an area, you can manually configure a metric value on each ABR. The metric assigns a priority to the ABR for its selection as the priority default route in an area. The following command is used for setting the metric value:

```
>> # /cfg/l3/ospf/default <metric value> <metric type (1 or 2)>
```

where *<metric value>* sets the priority for choosing this switch for default route. The value *none* sets no default and 1 sets the highest priority for default route. Metric type determines the method for influencing routing decisions for external routes.

To clear a default route metric from the switch, use the following command:

```
>> # /cfg/l3/ospf/default none
```

Virtual links

Usually, all areas in an OSPF AS are physically connected to the backbone. In some cases where this is not possible, you can use a virtual link. Virtual links are created to connect one area to the backbone through another non-backbone area.

The area which contains a virtual link must be a transit area and have full routing information. Virtual links cannot be configured inside a stub area or NSSA. The area type must be defined as *transit* using the following command:

```
>> # /cfg/l3/ospf/aindex <area index>/type transit
```

The virtual link must be configured on the routing devices at each endpoint of the virtual link, though they may traverse multiple routing devices. To configure a switch as one endpoint of a virtual link, use the following command:

```
>> # /cfg/l3/ospf/virt <link number>/aindex <area index>/nbr <router ID>
```

where *<link number>* is a value between 1 and 3, *<area index>* is the OSPF area index of the transit area, and *<router ID>* is the IP address of the virtual neighbor (*nbr*), the routing device at the target endpoint. Another router ID is needed when configuring a virtual link in the other direction. To provide the switch with a router ID, see “Router ID.”

For a detailed configuration example on Virtual Links, see “Example 2: Virtual Links.”

Router ID

Routing devices in OSPF areas are identified by a router ID. The router ID is expressed in IP address format. The IP address of the router ID is not required to be included in any IP interface range or in any OSPF area.

The router ID can be configured in one of the following two ways:

- Dynamically—OSPF protocol configures the lowest IP interface IP address as the router ID. This is the default.
- Statically—Use the following command to manually configure the router ID

```
>> # /cfg/l3/rtrid <IP address>
```

To modify the router ID from static to dynamic, set the router ID to 0.0.0.0, save the configuration, and reboot the switch. To view the router ID, enter:

```
>> # /info/l3/ospf/gen
```

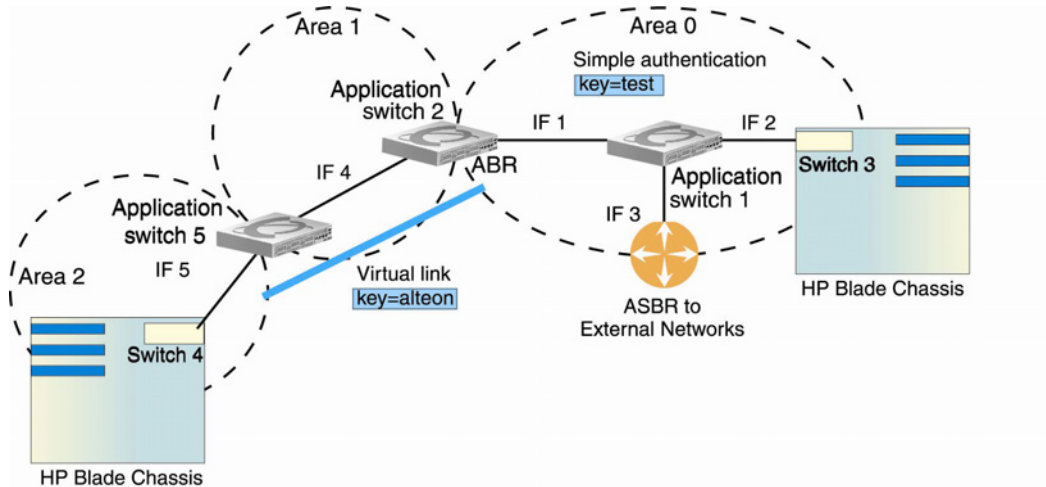
Authentication

OSPF protocol exchanges can be authenticated so that only trusted routing devices can participate. This ensures less processing on routing devices that are not listening to OSPF packets.

OSPF allows packet authentication and uses IP multicast when sending and receiving packets. Routers participate in routing domains based on predefined passwords. The switch software supports simple password (type 1 plain text passwords) and MD5 cryptographic authentication. This type of authentication allows a password to be configured per area.

The following figure shows authentication configured for area 0 with the password test. Simple authentication is also configured for the virtual link between area 2 and area 0. Area 1 is not configured for OSPF authentication.

Figure 20 OSPF authentication



To configure simple plain text OSPF passwords on the switches shown in the figure use the following commands:

1. Enable OSPF authentication for Area 0 on switches 1, 2, and 3.

```
>> # /cfg/l3/ospf/aindex 0/auth password
```

2. Configure a simple text password up to eight characters for each OSPF IP interface in Area 0 on switches 1, 2, and 3.

```
>> # /cfg/l3/ospf/if 1
>> OSPF Interface 1 # key test
>> OSPF Interface 1 # ../if 2
>> OSPF Interface 2 # key test
>> OSPF Interface 1 # ../if 3
>> OSPF Interface 3 # key test
```

3. Enable OSPF authentication for Area 2 on switch 4.

```
>> # /cfg/l3/ospf/aindex 2/auth password
```

4. Configure a simple text password up to eight characters for the virtual link between Area 2 and Area 0 on switches 2 and 4.

```
>> # /cfg/l3/ospf/virt 1/key packard
```

Use the following commands to configure MD5 authentication on the switches shown in the figure:

1. Enable OSPF MD5 authentication for Area 0 on switches 1, 2, and 3

```
>> # /cfg/l3/ospf/aindex 0/auth md5
```

2. Configure MD5 key ID for Area 0 on switches 1, 2, and 3.

```
>> # /cfg/l3/ospf/md5key 1/key test
```

3. Assign MD5 key ID to OSPF interfaces on switches 1, 2, and 3.

```
>> # /cfg/l3/ospf/if 1
>> OSPF Interface 1 # mdkey 1
>> OSPF Interface 1 # ../if 2
>> OSPF Interface 2 # mdkey 1
>> OSPF Interface 1 # ../if 3
>> OSPF Interface 3 # mdkey 1
```

4. Enable OSPF MD5 authentication for Area 2 on switch 4.

```
>> # /cfg/l3/ospf/aindex 2/auth md5
```

5. Configure MD5 key for the virtual link between Area 2 and Area 0 on switches 2 and 4.

```
>> # /cfg/l3/ospf/md5key 2/key packard
```

6. Assign MD5 key ID to OSPF virtual link on switches 2 and 4.

```
>> # /cfg/l3/ospf/virt 1/mdkey 2
```

Host routes for load balancing

The HP 10GbE switch implementation of OSPF includes host routes. Host routes are used for advertising network device IP addresses to external networks, accomplishing the following goals:

- **ABR Load Sharing**
As a form of load balancing, host routes can be used for dividing OSPF traffic among multiple ABRs. To accomplish this, each switch provides identical services but advertises a host route for a different IP address to the external network. If each IP address serves a different and equal portion of the external world, incoming traffic from the upstream router should be split evenly among ABRs.
- **ABR Failover**
Complementing ABR load sharing, identical host routes can be configured on each ABR. These host routes can be given different costs so that a different ABR is selected as the preferred route for each server and the others are available as backups for failover purposes.
- **Equal Cost Multipath (ECMP)**
With equal cost multipath, a router potentially has several available next hops towards any given destination. ECMP allows separate routes to be calculated for each IP Type of Service. All paths of equal cost to a given destination are calculated, and the next hops for all equal-cost paths are inserted into the routing table.

If redundant routes via multiple routing processes (such as OSPF, RIP, BGP, or static routes) exist on your network, the switch defaults to the OSPF-derived route.

OSPF features not supported in this release

The following OSPF features are not supported in this release:

- Summarizing external routes
- Filtering OSPF routes
- Using OSPF to forward multicast routes
- Configuring OSPF on non-broadcast multi-access networks (such as frame relay, X.25, and ATM)

OSPF configuration examples


A summary of the basic steps for configuring OSPF on the switch is listed here. Detailed instructions for each of the steps are covered in the following sections:

- Configure IP interfaces.
One IP interface is required for each desired network (range of IP addresses) being assigned to an OSPF area on the switch.
- (Optional) Configure the router ID.
The router ID is required only when configuring virtual links on the switch.
- Enable OSPF on the switch.
- Define the OSPF areas.
- Configure OSPF interface parameters.
IP interfaces are used for attaching networks to the various areas.
- (Optional) Configure route summarization between OSPF areas.
- (Optional) Configure virtual links.
- (Optional) Configure host routes.

Example 1: Simple OSPF domain (CLI example)

In this example, two OSPF areas are defined—one area is the backbone and the other is a stub area. A stub area does not allow advertisements of external routes, thus reducing the size of the database. Instead, a default summary route of IP address 0.0.0.0 is automatically inserted into the stub area. Any traffic for IP address destinations outside the stub area will be forwarded to the stub area's IP interface, and then into the backbone.

Figure 21 Simple OSPF domain



Follow this procedure to configure OSPF support as shown in the figure.

1. Configure IP interfaces on each network that will be attached to OSPF areas.
2. In this example, two IP interfaces are needed: one for the backbone network on 10.10.7.0/24 and one for the stub area network on 10.10.12.0/24.

```
>> # /cfg/l3/if 1 (Select menu for IP interface 1)
>> IP Interface 1 # addr 10.10.7.1(Set IP address on backbone network)
>> IP Interface 1 # mask 255.255.255.0(Set IP mask on backbone network)
>> IP Interface 1 # enable(Enable IP interface 1)
>> IP Interface 1 # ../if 2(Select menu for IP interface 2)
>> IP Interface 2 # addr 10.10.12.1(Set IP address on stub area network)
>> IP Interface 2 # mask 255.255.255.0(Set IP mask on stub area network)
>> IP Interface 2 # enable(Enable IP interface 2)
```

3. Enable OSPF.

```
>> IP Interface 2 # /cfg/l3/ospf/on(Enable OSPF on the switch)
```

4. Define the backbone.

5. The backbone is always configured as a transit area using areaid 0.0.0.0.

```
>> Open Shortest Path First # aindex 0(Select menu for area index 0)
>> OSPF Area (index) 0 # areaid 0.0.0.0(Set the ID for backbone area 0)
>> OSPF Area (index) 0 # type transit(Define backbone as transit type)
>> OSPF Area (index) 0 # enable(Enable the area)
```

6. Define the stub area.

```
>> OSPF Area (index) 0 # ../aindex 1(Select menu for area index 1)
>> OSPF Area (index) 1 # areaid 0.0.0.1(Set the area ID for OSPF area 1)
>> OSPF Area (index) 1 # type stub(Define area as stub type)
>> OSPF Area (index) 1 # enable(Enable the area)
```

7. Attach the network interface to the backbone.

```
>> OSPF Area 1 # ../if 1(Select OSPF menu for IP interface 1)
>> OSPF Interface 1 # aindex 0(Attach network to backbone index)
>> OSPF Interface 1 # enable(Enable the backbone interface)
```

8. Attach the network interface to the stub area.

```
>> OSPF Interface 1 # ../if 2(Select OSPF menu for IP interface 2)
>> OSPF Interface 2 # aindex 1(Attach network to stub area index)
>> OSPF Interface 2 # enable(Enable the stub area interface)
```

9. Apply and save the configuration changes.

```
>> OSPF Interface 2 # apply (Global command to apply all changes)
>> OSPF Interface 2 # save (Global command to save all changes)
```

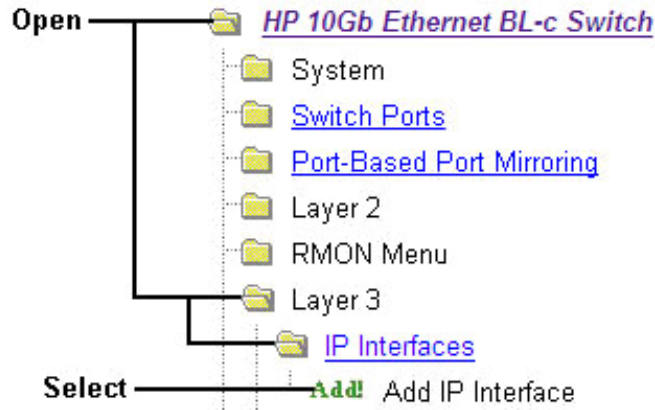
Example 1: Simple OSPF domain (BBI example)

1. Configure IP interfaces on each network that will be attached to OSPF areas:

- IF 1
 - IP address = 10.10.7.1
 - Subnet mask = 255.255.255.0
- IF 2
 - IP address = 10.10.12.1
 - Subnet mask = 255.255.255.0

- a. Click the **Configure** context button.

- b. Open the IP Interfaces folder, and select **Add IP Interface**.

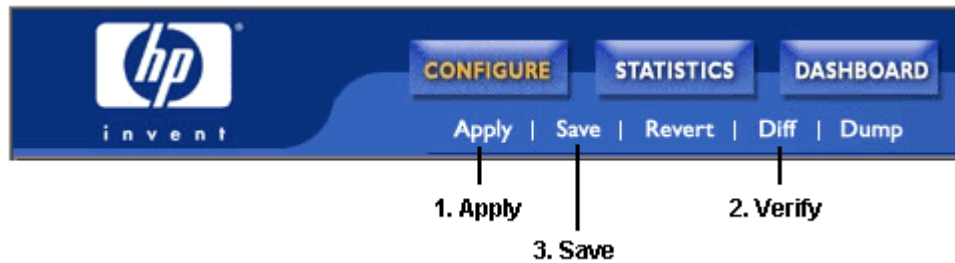


- c. Configure an IP interface. Enter the IP address, subnet mask, and enable the interface.


IP Interface Configuration

IP Interface Identifier (1-250)	<input type="text" value="1"/>
IP Address	<input type="text" value="10.10.7.1"/>
Enabled?	Enabled <input type="button" value="v"/>
Subnet Mask	<input type="text" value="255.255.255.0"/>
VLAN Membership ID (1 - 4095)	<input type="text" value="1"/>
Enable/Disable BOOTP Relay	Enabled <input type="button" value="v"/>

- d. Click **Submit**.
2. Apply, verify, and save the configuration.



3. Enable OSPF.
 - a. Open the OSPF Routing Protocol folder, and select **General**.



- b. Enable OSPF.

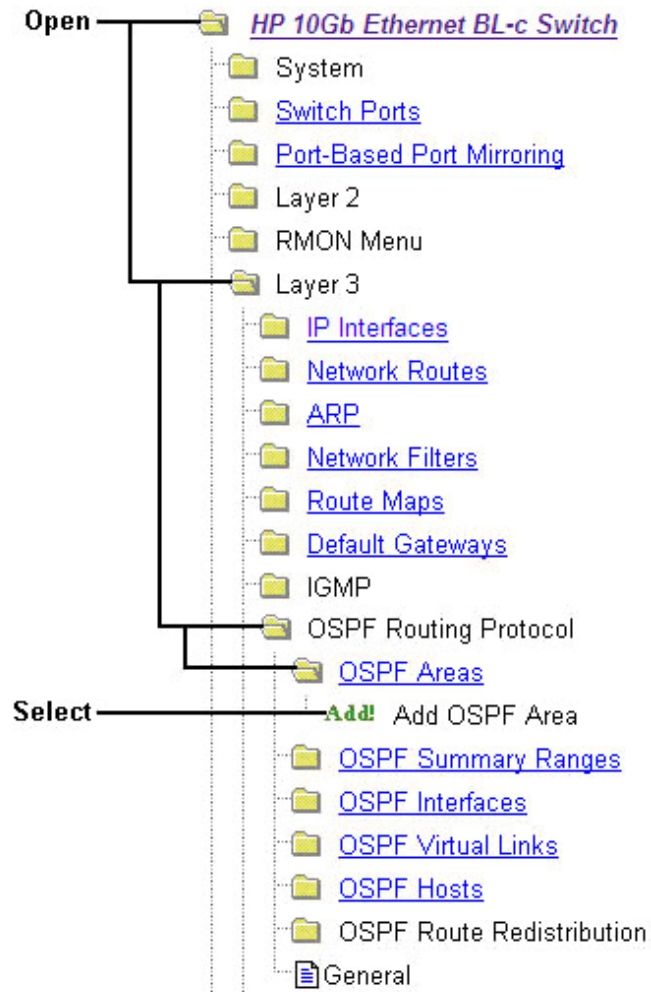
OSPF General Configuration

Globally Enable OSPF ?	<input type="text" value="enabled"/>
External LSDB Limit (0-2000)	<input type="text" value="0"/>
Default Route Metric (1-16777215, 0=none)	<input type="text" value="0"/>
Default Route Metric Type	<input type="text" value="none"/>

OSPF MD5 Keys Configuration

ID MD5 Key

- c. Click **Submit**.
4. Configure OSPF Areas.
 - a. Open the OSPF Areas folder, and select **Add OSPF Area**.



- b. Configure the OSPF backbone area 0.

OSPF Area Configuration

Area Number (0-2)	<input type="text" value="0"/>
Area ID	<input type="text" value="0.0.0.0"/>
Enabled?	<input type="text" value="enabled"/>
Area Type	<input type="text" value="transit"/>
Stub Area Metric (1-65535)	<input type="text" value="1"/>
SPF Interval (0-255)	<input type="text" value="10"/>
Authentication Type?	<input type="text" value="none"/>

- c. Click **Submit**.
- d. Select **Add OSPF Area**.
- e. Configure the OSPF area 1.

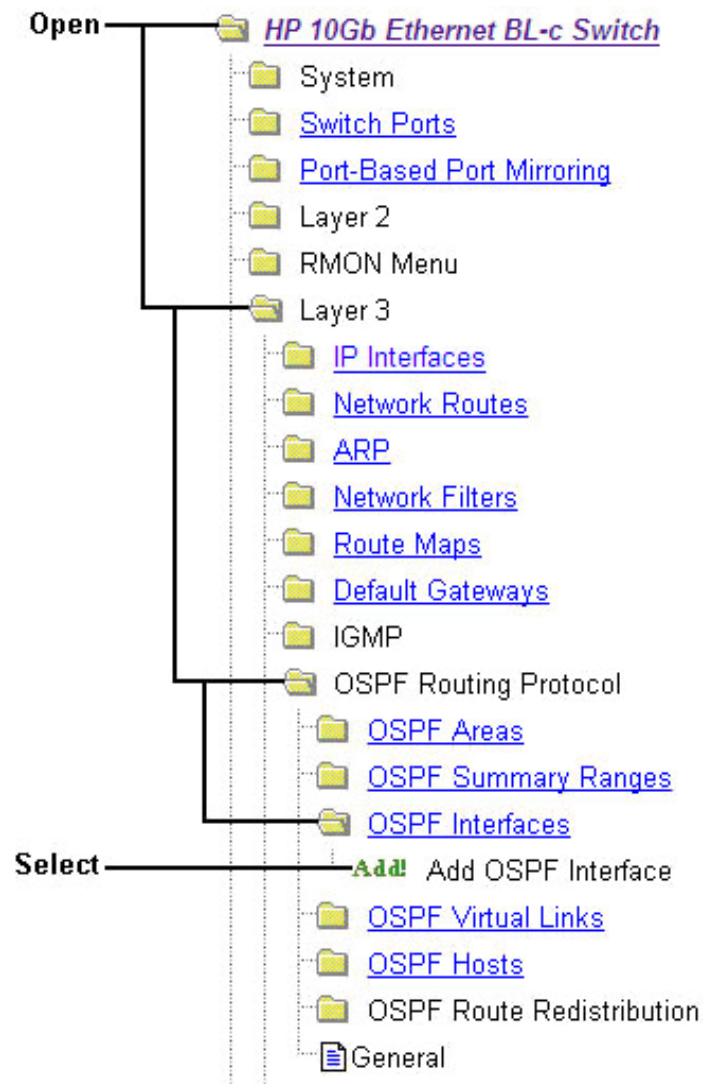
OSPF Area Configuration

Area Number (0-2)	<input type="text" value="1"/>
Area ID	<input type="text" value="0.0.0.1"/>
Enabled?	<input type="text" value="enabled"/>
Area Type	<input type="text" value="stub"/>
Stub Area Metric (1-65535)	<input type="text" value="1"/>
SPF Interval (0-255)	<input type="text" value="10"/>
Authentication Type?	<input type="text" value="none"/>

- f. Click **Submit**.

5. Configure OSPF Interfaces.

- a. Open the OSPF Interfaces folder, and select **Add OSPF Interface**.

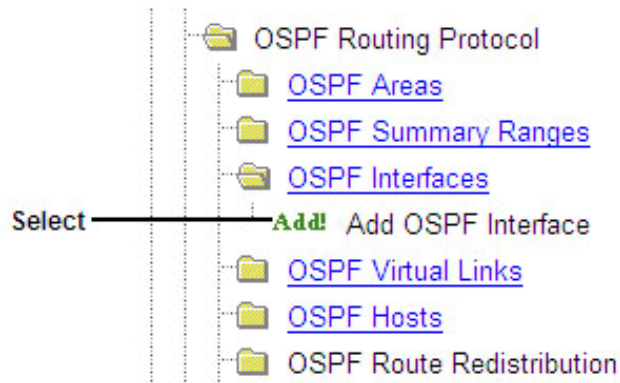


- b. Configure the OSPF Interface 1, and attach it to the backbone area 0.

OSPF Interface Configuration

IP Interface Identifier (1-255)	<input type="text" value="1"/>
Area Number (0-2)	<input type="text" value="0"/>
Enabled?	<input type="button" value="enabled"/>
Router Priority (0-255)	<input type="text" value="1"/>
Output Cost (1-65535)	<input type="text" value="1"/>
Hello Interval (1-65535 sec)	<input type="text" value="10"/>
Dead Interval (1-65535 sec)	<input type="text" value="40"/>
Transit Delay (1-3600 sec)	<input type="text" value="1"/>
Retransmit Interval (1-3600 sec)	<input type="text" value="5"/>
Authentication Key	<input type="text"/>
MD5 Key ID (1-255)	<input type="text" value="0"/>

- c. Click **Submit**.
- d. Select **Add OSPF Interface**.



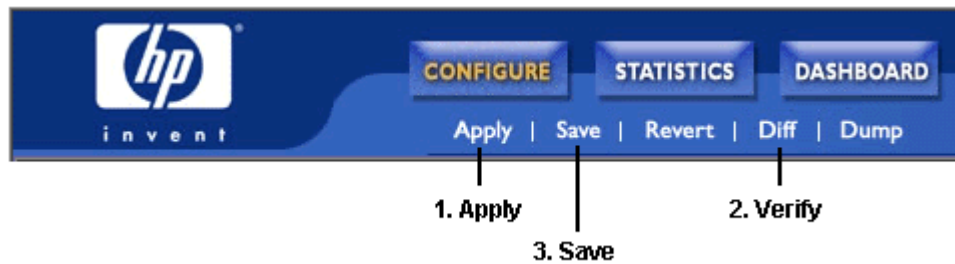
- e. Configure the OSPF Interface 2, and attach it to the stub area 1.

OSPF Interface Configuration

IP Interface Identifier (1-255)	<input type="text" value="2"/>
Area Number (0-2)	<input type="text" value="1"/>
Enabled?	<input type="button" value="enabled"/>
Router Priority (0-255)	<input type="text" value="1"/>
Output Cost (1-65535)	<input type="text" value="1"/>
Hello Interval (1-65535 sec)	<input type="text" value="10"/>
Dead Interval (1-65535 sec)	<input type="text" value="40"/>
Transit Delay (1-3600 sec)	<input type="text" value="1"/>
Retransmit Interval (1-3600 sec)	<input type="text" value="5"/>
Authentication Key	<input type="text"/>
MD5 Key ID (1-255)	<input type="text" value="0"/>

- f. Click **Submit**.


6. Apply, verify, and save the configuration.



Example 2: Virtual links

In the example shown in the following figure, area 2 is not physically connected to the backbone as is usually required. Instead, area 2 will be connected to the backbone via a virtual link through area 1. The virtual link must be configured at each endpoint.

Figure 22 Configuring a virtual link



Configuring OSPF for a virtual link on Switch A

1. Configure IP interfaces on each network that will be attached to the switch.

In this example, two IP interfaces are needed on Switch A: one for the backbone network on 10.10.7.0/24 and one for the transit area network on 10.10.12.0/24.

```
>> # /cfg/l3/if 1 (Select menu for IP interface 1)
>> IP Interface 1 # addr 10.10.7.1 (Set IP address on backbone network)
>> IP Interface 1 # mask 255.255.255.0 (Set IP mask on backbone network)
>> IP Interface 1 # enable (Enable IP interface 1)
>> IP Interface 1 # ../if 2 (Select menu for IP interface 2)
>> IP Interface 2 # addr 10.10.12.1 (Set IP address on
transit area network)
>> IP Interface 2 # mask 255.255.255.0 (Set IP mask on
transit area network)
>> IP Interface 2 # enable (Enable interface 2)
```

2. Configure the router ID.
3. A router ID is required when configuring virtual links. Later, when configuring the other end of the virtual link on Switch B, the router ID specified here will be used as the target virtual neighbor (nbr) address

```
>> IP Interface 1 # /cfg/l3/rtrid 10.10.10.1 (Set static router ID)
```

4. Enable OSPF.

```
>> IP # /cfg/l3/ospf/on
```

5. Define the backbone.

```
>> Open Shortest Path First # aindex 0 (Select menu for area index 0)
>> OSPF Area (index) 0 # areaid 0.0.0.0 (Set the area ID)
>> OSPF Area (index) 0 # type transit (Define backbone as transit type)
>> OSPF Area (index) 0 # enable (Enable the area)
```

6. Define the transit area.

7. The area that contains the virtual link must be configured as a transit area

```
>> OSPF Area (index) 0 # ../aindex 1 (Select menu for area index 1)
>> OSPF Area (index) 1 # areaid 0.0.0.1 (Set the area ID for OSPF area 1)
>> OSPF Area (index) 1 # type transit (Define area as transit type)
>> OSPF Area (index) 1 # enable (Enable the area)
```

8. Attach the network interface to the backbone.

```
>> OSPF Area (index) 1 # ../if 1 (Select OSPF menu for IP interface 1)
>> OSPF Interface 1 # aindex 0 (Attach network to backbone index)
>> OSPF Interface 1 # enable (Enable the backbone interface)
```

9. Attach the network interface to the transit area.

```
>> OSPF Interface 1 # ../if 2 (Select OSPF menu for IP interface 2)
>> OSPF Interface 2 # aindex 1 (Attach network to transit area index)
>> OSPF Interface 2 # enable (Enable the transit area interface)
```

10. Configure the virtual link.

11. The `nbr` router ID configured in this step must be the same as the router ID that will be configured for Switch B in step 2.

```
>> OSPF Interface 2 # ../virt 1 (Specify a virtual link number)
>> OSPF Virtual Link 1 # aindex 1 (Specify the transit area
for the virtual link)
>> OSPF Virtual Link 1 # nbr 10.10.14.1 (Specify the router ID
of the recipient)
>> OSPF Virtual Link 1 # enable (Enable the virtual link)
```

12. Apply and save the configuration changes

```
>> OSPF Interface 2 # apply (Apply all changes)
>> OSPF Interface 2 # save (Save all changes)
```

Configuring OSPF for a virtual link on Switch B

1. Configure IP interfaces on each network that will be attached to OSPF areas.
2. Two IP interfaces are needed on Switch B: one for the transit area network on 10.10.12.0/24 and one for the stub area network on 10.10.24.0/24.

```
>> # /cfg/l3/if 1 (Select menu for IP interface 1)
>> IP Interface 1 # addr 10.10.12.2 (Set IP address
on transit area network)
>> IP Interface 1 # mask 255.255.255.0 (Set IP mask
on transit area network)
>> IP Interface 1 # enable (Enable IP interface 1)
>> IP Interface 1 # ../if 2 (Select menu for IP interface 2)
>> IP Interface 2 # addr 10.10.24.1 (Set IP address on stub area network)
>> IP Interface 2 # mask 255.255.255.0 (Set IP mask on stub area network)
>> IP Interface 2 # enable (Enable IP interface 2)
```

3. Configure the router ID.

4. A router ID is required when configuring virtual links. This router ID should be the same one specified as the target virtual neighbor (`nbr`) on Switch A.

```
>> IP Interface 2 # /cfg/l3/rtrid 10.10.14.1 (Set static router ID)
```

5. Enable OSPF.

```
>> IP# /cfg/l3/ospf/on
```

6. Define the backbone.

7. This version of switch software requires that a backbone index be configured on the non-backbone end of the virtual link as follows:

```
>> Open Shortest Path First # aindex 0 (Select the menu for area index 0)
>> OSPF Area (index) 0 # areaid 0.0.0.0 (Set the area ID for OSPF area 0)
>> OSPF Area (index) 0 # enable (Enable the area)
```

8. Define the transit area.

```
>> OSPF Area (index) 0 # ../aindex 1 (Select menu for area index 1)
>> OSPF Area (index) 1 # areaid 0.0.0.1 (Set the area ID for OSPF area 1)
>> OSPF Area (index) 1 # type transit (Define area as transit type)
>> OSPF Area (index) 1 # enable (Enable the area)
```

9. Define the stub area.

```
>> OSPF Area (index) 1 # ../aindex 2 (Select the menu for area index 2)
>> OSPF Area (index) 2 # areaid 0.0.0.2 (Set the area ID for OSPF area 2)
>> OSPF Area (index) 2 # type stub (Define area as stub type)
>> OSPF Area (index) 2 # enable (Enable the area)
```

10. Attach the network interface to the backbone.

```
>> OSPF Area (index) 2 # ../if 1 (Select OSPF menu for IP interface 1)
>> OSPF Interface 1 # aindex 1 (Attach network to transit area index)
>> OSPF Interface 1 # enable (Enable the transit area interface)
```

11. Attach the network interface to the transit area.

```
>> OSPF Interface 1 # ../if 2 (Select OSPF menu for IP interface 2)
>> OSPF Interface 2 # aindex 2 (Attach network to stub area index)
>> OSPF Interface 2 # enable (Enable the stub area interface)
```

12. Configure the virtual link.**13. The nbr router ID configured in this step must be the same as the router ID that was configured for Switch A in step 2.**

```
>> OSPF Interface 2 # ../virt 1 (Specify a virtual link number)
>> OSPF Virtual Link 1 # aindex 1 (Specify the transit area
for the virtual link)
>> OSPF Virtual Link 1 # nbr 10.10.10.1 (Specify the router ID
of the recipient)
>> OSPF Virtual Link 1 # enable (Enable the virtual link)
```

14. Apply and save the configuration changes.

```
>> OSPF Interface 2 # apply (Apply all changes)
>> OSPF Interface 2 # save (Save all changes)
```

Other Virtual Link Options

- You can use redundant paths by configuring multiple virtual links.
- Only the endpoints of the virtual link are configured. The virtual link path may traverse multiple routers in an area as long as there is a routable path between the endpoints.


Example 3: Summarizing routes

By default, ABRs advertise all the network addresses from one area into another area. Route summarization can be used for consolidating advertised addresses and reducing the perceived complexity of the network.

If the network IP addresses in an area are assigned to a contiguous subnet range, you can configure the ABR to advertise a single summary route that includes all the individual IP addresses within the area.

The following example shows one summary route from area 1 (stub area) injected into area 0 (the backbone). The summary route consists of all IP addresses from 36.128.192.0 through 36.128.254.255 except for the routes in the range 36.128.200.0 through 36.128.200.255.

Figure 23 Summarizing routes



NOTE: You can specify a range of addresses to prevent advertising by using the `hide` option. In this example, routes in the range 36.128.200.0 through 36.128.200.255 are kept private.

Follow this procedure to configure OSPF support on Switch A and Switch B, as shown in the figure.

1. Configure IP interfaces for each network which will be attached to OSPF areas.

```
>> # /cfg/l3/if 1(Select menu for IP interface 1)
>> IP Interface 1 # addr 10.10.7.1(Set IP address on backbone network)
>> IP Interface 1 # mask 255.255.255.0(Set IP mask on backbone network)
>> IP Interface 1 # ena(Enable IP interface 1)
>> IP Interface 1 # ../if 2(Select menu for IP interface 2)
>> IP Interface 2 # addr 36.128.192.1(Set IP address on stub area network)
>> IP Interface 2 # mask 255.255.192.0(Set IP mask on stub area network)
>> IP Interface 2 # ena(Enable IP interface 2)
```

2. Enable OSPF.

```
>> IP Interface 2 # /cfg/l3/ospf/on
```

3. Define the backbone.

```
>> Open Shortest Path First # aindex 0 (Select menu for area index 0)
>> OSPF Area (index) 0 # areaid 0.0.0.0(Set the ID for backbone area 0)
>> OSPF Area (index) 0 # type transit (Define backbone as transit type)
>> OSPF Area (index) 0 # enable (Enable the area)
```

4. Define the stub area.

```
>> OSPF Area (index) 0 # ../aindex 1 (Select menu for area index 1)
>> OSPF Area (index) 1 # areaid 0.0.0.1 (Set the area ID for OSPF
area 1)
>> OSPF Area (index) 1 # type stub (Define area as stub type)
>> OSPF Area (index) 1 # enable (Enable the area)
```

5. Attach the network interface to the backbone.

```
>> OSPF Area (index) 1 # ../if 1 (Select OSPF menu for IP interface 1)
>> OSPF Interface 1 # aindex 0 (Attach network to backbone index)
>> OSPF Interface 1 # enable (Enable the backbone interface)
```

6. Attach the network interface to the stub area.

```
>> OSPF Interface 1 # ../if 2 (Select OSPF menu for IP interface 2)
>> OSPF Interface 2 # aindex 1 (Attach network to stub area index)
>> OSPF Interface 2 # enable (Enable the stub area interface)
```

7. Configure route summarization by specifying the starting address and mask of the range of addresses to be summarized.

```
>> OSPF Interface 2 # ../range 1 (Select menu for summary range)
>> OSPF Summary Range 1 # addr 36.128.192.0 (Set base IP address
                                         of summary range)
>> OSPF Summary Range 1 # mask 255.255.192.0 (Set mask address
                                         for summary range)
>> OSPF Summary Range 1 # aindex 0 (Inject summary route into backbone)
>> OSPF Summary Range 1 # enable (Enable summary range)
```

8. Use the hide command to prevent a range of addresses from advertising to the backbone.

```
>> OSPF Interface 2 # ../range 2 (Select menu for summary range)
>> OSPF Summary Range 2 # addr 36.128.200.0 (Set base IP address)
>> OSPF Summary Range 2 # mask 255.255.255.0 (Set mask address)
>> OSPF Summary Range 2 # hide enable (Hide the range of addresses)
```

9. Apply and save the configuration changes.

```
>> OSPF Summary Range 2 # apply (Apply all changes)
>> OSPF Summary Range 2 # save (Save all changes)
```

Verifying OSPF configuration

Use the following commands to verify the OSPF configuration on your switch:

- /info/13/ospf/general
- /info/13/ospf/nbr
- /info/13/ospf/dbase/dbsum
- /info/13/ospf/routes
- /stats/13/route

See the *HP 10Gb Ethernet BL-c Switch Command Reference* for information on the above commands.

Remote monitoring

Introduction

Remote Monitoring (RMON) allows network devices to exchange network monitoring data.

RMON performs the following major functions:

- Gathers cumulative statistics for Ethernet interfaces
- Tracks a history of statistics for Ethernet interfaces
- Creates and triggers alarms for user-defined events

Overview

The RMON MIB provides an interface between the RMON agent on the switch and an RMON management application. The RMON MIB is described in RFC 1757.

The RMON standard defines objects that are suitable for the management of Ethernet networks. The RMON agent continuously collects statistics and proactively monitors switch performance. RMON allows you to monitor traffic flowing through the switch.

The switch supports the following RMON Groups, as described in RFC 1757:

- Group 1: Statistics
- Group 2: History
- Group 3: Alarms
- Group 9: Events

RMON group 1 —statistics

The switch supports collection of Ethernet statistics as outlined in the RMON statistics MIB, in reference to `etherStatsTable`. You can enable RMON statistics on a per-port basis, and you can view them using the following command: `/stat/port x/rmon`. RMON statistics are sampled every second, and new data overwrites any old data on a given port.

NOTE: RMON port statistics must be enabled for the port before you can view RMON statistics.

Configuring RMON Statistics (CLI example)

1. Enable RMON on each port where you wish to collect RMON statistics.


```
>> /cfg/port 20/rmon (Select Port 20 RMON)
>> Port 20 RMON# ena (Enable RMON)
>> Port 20 RMON# apply (Make your changes active)
>> Port 20 RMON# save (Save for restore after reboot)
```

2. View RMON statistics for the port.

```
>> /stats/port 20 (Select Port 20 Stats)
>> Port Statistics# rmon
-----
RMON statistics for port 20:
etherStatsDropEvents: NA
etherStatsOctets: 7305626
etherStatsPkts: 48686
etherStatsBroadcastPkts: 4380
etherStatsMulticastPkts: 6612
etherStatsCRCAlignErrors: 22
etherStatsUndersizePkts: 0
etherStatsOversizePkts: 0
etherStatsFragments: 2
etherStatsJabbers: 0
etherStatsCollisions: 0
etherStatsPkts64Octets: 27445
etherStatsPkts65to127Octets: 12253
etherStatsPkts128to255Octets: 1046
etherStatsPkts256to511Octets: 619
etherStatsPkts512to1023Octets: 7283
etherStatsPkts1024to1518Octets: 38
```

Configuring RMON Statistics (BBI example)

1. Configure ports.
 - a. Click the **Configure** context button.
 - b. Select **Switch Ports** (click the underlined text, not the folder).



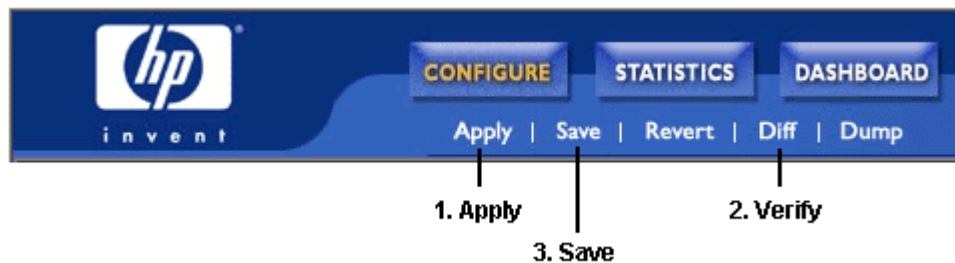
2. Select a port.

Switch Ports Configuration								
Switch Port	State	VLAN Tagging	Default PVID	PVID tagging	Multicast Threshold	Broadcast Threshold	Destination Lookup Fail Threshold	802.1p Priority
1	enabled	disabled	1	enabled	disabled	disabled	disabled	0
2	enabled	disabled	1	enabled	disabled	disabled	disabled	0
3	enabled	disabled	1	enabled	disabled	disabled	disabled	0
4	enabled	disabled	1	enabled	disabled	disabled	disabled	0
⋮								
18	enabled	disabled	1	enabled	disabled	disabled	disabled	0
19	enabled	disabled	1	enabled	disabled	disabled	disabled	0
Select → 20	enabled	disabled	1	enabled	disabled	disabled	disabled	0
21	enabled	disabled	1	enabled	disabled	disabled	disabled	0

3. Enable RMON on the port.

Switch Port 20 Configuration	
Switch Port State	Enabled ▾
RMON Instrumentation	Enabled ▾
VLAN Tagging	Enabled ▾
PVID Tagging	Disabled ▾
Port STP	On ▾
Default Port VLAN ID (1 - 4095)	1
Flow Control	both Rx/Tx ▾
Autonegotiation	Off
Speed	10000
Duplex Mode	Full
Enable/Disable sending Link UP/Down Trap	Enabled ▾
Port Name	Uplink3

4. Click **Submit**.
5. Apply, verify, and save the configuration.



RMON group 2—history

The RMON History group allows you to sample and archive Ethernet statistics for a specific interface during a specific time interval.

NOTE: RMON port statistics must be enabled for the port before an RMON history group can monitor the port.

Data is stored in buckets, which store data gathered during discreet sampling intervals. At each configured interval, the history instance takes a sample of the current Ethernet statistics, and places them into a bucket. History data buckets reside in dynamic memory. When the switch is re-booted, the buckets are emptied.

Requested buckets (`/cfg/rmon/hist x/rbnum`) are the number of buckets, or data slots, requested by the user for each History Group. Granted buckets (`/info/rmon/hist x/gbnum`) are the number of buckets granted by the system, based on the amount of system memory available. The system grants a maximum of 50 buckets.

Use an SNMP browser to view History samples.

History MIB objects

The type of data that can be sampled must be of an ifIndex object type, as described in RFC1213 and RFC1573. The most common data type for the history sample is as follows:

```
1.3.6.1.2.1.2.2.1.1.x -mgmt.interfaces.ifTable.ifIndex.interface
```

The last digit (*x*) represents the interface on which to monitor, which corresponds to the port number (1-16, 18-21). History sampling is done per port, by utilizing the interface number to specify the port number.

Configure RMON History (CLI example)

1. Enable RMON on each port where you wish to collect RMON History.

```
>> /cfg/port 21/rmon                (Select Port 21 RMON)
>> Port 21# ena                    (Enable RMON)
>> Port 21 RMON# apply              (Make your changes active)
>> Port 21 RMON# save               (Save for restore after reboot)
```

2. Configure the RMON History parameters.

```
>> /cfg/rmon/hist 1                  (Select RMON History 1)
>> RMON History 1# ifoid 1.3.6.1.2.1.2.2.1.1.21
>> RMON History 1# rbnum 30
>> RMON History 1# intrval 120
>> RMON History 1# owner "Owner_History_1"
```


3. Apply and save the configuration.

```
>> RMON History 1# apply              (Make your changes active)
>> RMON History 1# save               (Save for restore after reboot)
```

This configuration creates an RMON History group to monitor port 21. It takes a data sample every two minutes, and places the data into one of the 30 requested buckets. After 30 samples are gathered, the new samples overwrite the previous samples, beginning with the first bucket. Use SNMP to view the data.

Configure RMON History (BBI example)

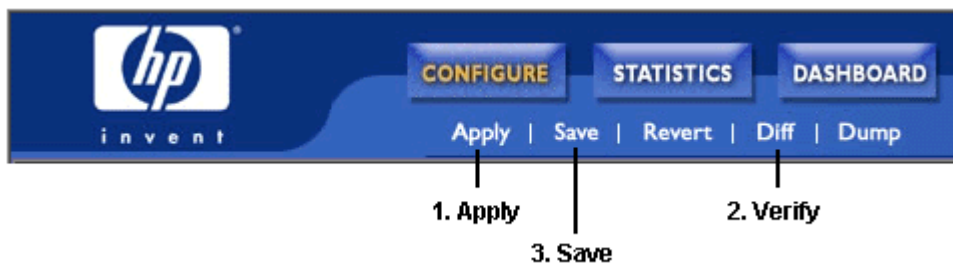
1. Configure an RMON History group.
 - a. Click the **Configure** context button.
 - b. Open the Switch folder, and select **RMON > History > Add History Group**.



2. Configure RMON History Group parameters.

RMON History Configuration	
History Group ID (1 - 65535)	1
MIB Object ID	1.3.6.2.1.2.2.1.1.21
Number of Buckets Requested (1 - 65535)	30
Polling Interval (1 - 3600)	120
Owner	Owner_History_1
<input type="button" value="Submit"/> <input type="button" value="Delete"/>	

3. Click **Submit**.
4. Apply, verify, and save the configuration.



RMON group 3—alarms

The RMON Alarm group allows you to define a set of thresholds used to determine network performance. When a configured threshold is crossed, an alarm is generated. For example, you can configure the switch to issue an alarm if more than 1,000 CRC errors occur during a 10-minute time interval.

Each Alarm index consists of a variable to monitor, a sampling time interval, and parameters for rising and falling thresholds. The Alarm group can be used to track rising or falling values for a MIB object. The object must be a counter, gauge, integer, or time interval.

Use the `/cfg/rmon/alarm x/revtidx` or `/fevtidx` to correlate an alarm index to an event index. When the alarm threshold is reached, the corresponding event is triggered.

Alarm MIB objects

The most common data types used for alarm monitoring are ifStats: errors, drops, bad CRCs, and so on. These MIB Object Identifiers (OIDs) correlate to the ones tracked by the History group. An example of an ICMP stat is as follows:

```
1.3.6.1.2.1.5.1.0 - mgmt.icmp.icmpInMsgs
```

The last digit (*x*) represents the interface on which to monitor, which corresponds to the interface number, or port number, as follows:

- 1-250 = IF 1-250
- 251 = port 1
- 252 = port 2
- 271 = port 21

This value represents the alarm's MIB OID, as a string. Note that for non-tables, you must supply a `.0` to specify end node.

Configure RMON Alarms (CLI example 1)

1. Configure the RMON Alarm parameters to track the number of packets received on a port.

```
>> /cfg/rmon/alarm 6 (Select RMON Alarm 6)
>> RMON Alarm 6# oid 1.3.6.1.2.1.2.2.1.10.270
>> RMON Alarm 6# intrval 3600
>> RMON Alarm 6# almttype rising
>> RMON Alarm 6# rlimit 2000000000
>> RMON Alarm 6# revtidx 6
>> RMON Alarm 6# sample abs
>> RMON Alarm 6# owner "Alarm_for_ifInOctets"
```

2. Apply and save the configuration.

```
>> RMON Alarm 6# apply (Make your changes active)
>> RMON Alarm 6# save (Save for restore after reboot)
```

This configuration creates an RMON alarm that checks `ifInOctets` on port 20 once every hour. If the statistic exceeds two billion, an alarm is generated that triggers event index 6.

Configure RMON Alarms (CLI example 2)

1. Configure the RMON Alarm parameters to track ICMP messages.

```
>> /cfg/rmon/alarm 5                                     (Select RMON Alarm 5)
>> RMON Alarm 5# oid 1.3.6.1.2.1.5.8.0
>> RMON Alarm 5# intrval 60
>> RMON Alarm 5# almtypе rising
>> RMON Alarm 5# rlimit 200
>> RMON Alarm 5# revtidx 5
>> RMON Alarm 5# sample delta
>> RMON Alarm 5# owner "Alarm_for_icmpInEchos"
```


2. Apply and save the configuration.

```
>> RMON Alarm 5# apply                                  (Make your changes active)
>> RMON Alarm 5# save                                  (Save for restore after reboot)
```

This configuration creates an RMON alarm that checks `icmpInEchos` on the switch once every minute. If the statistic exceeds 200 within a 60 second interval, an alarm is generated that triggers event index 5.

Configure RMON Alarms (BBI example 1)


1. Configure an RMON Alarm group.
 - a. Click the **Configure** context button.
 - b. Open the Switch folder, and select **RMON > Alarm > Add Alarm Group**.



- c. Configure RMON Alarm Group parameters to check `ifInOctets` on port 20 once every hour. Enter a rising limit of two billion, and a rising event index of 6. This configuration creates an RMON alarm that checks `ifInOctets` on port 20 once every hour. If the statistic exceeds two billion, an alarm is generated that triggers event index 6.

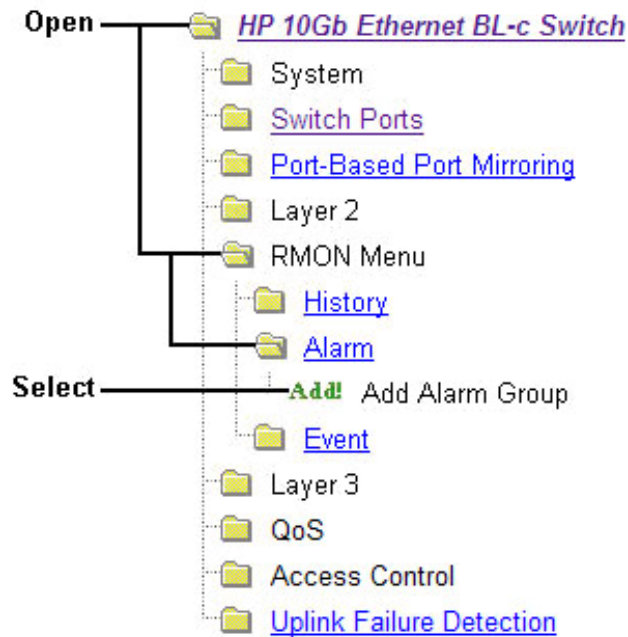
RMON Alarm Configuration	
Alarm Group ID (1 - 65535)	6
MIB Object ID	1.3.6.1.2.1.2.2.1.10.270
Rising Limit (-2147483647 - 2147483647)	2000000000
Falling Limit (-2147483647 - 2147483647)	0
Rising Event Index (0 - 65535)	6
Falling Event Index (0 - 65535)	0
Alarm Type	Rising ▾
Sample Type	Absolute ▾
Polling Interval (1 - 65535)	3600
Owner	Alarm_for_ifInOctets

2. Click **Submit**.
3. Apply, verify, and save the configuration.



Configure RMON Alarms (BBI example 2)

1. Configure an RMON Alarm group.
 - a. Click the **Configure** context button.
 - b. Open the Switch folder, and select **RMON > Alarm > Add Alarm Group**.

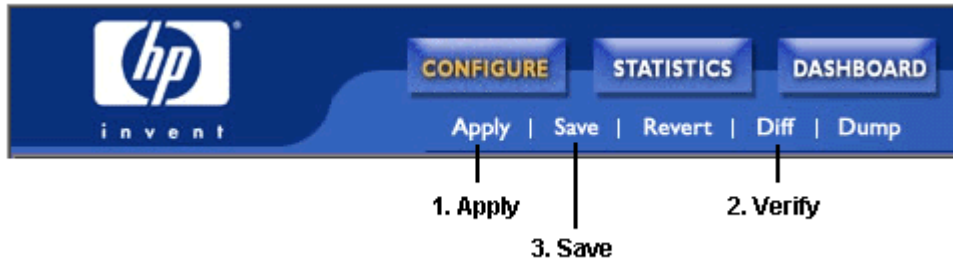


- c. Configure RMON Alarm Group parameters to check `icmpInEchos`, with a polling interval of 60, a rising limit of 200, and a rising event index of 5. This configuration creates an RMON alarm that checks `icmpInEchos` on the switch once every minute. If the statistic exceeds 200 within a 60 second interval, an alarm is generated that triggers event index 5.

RMON Alarm Configuration	
Alarm Group ID (1 - 65535)	5
MIB Object ID	1.3.6.1.2.1.5.8.0
Rising Limit (-2147483647 - 2147483647)	200
Falling Limit (-2147483647 - 2147483647)	0
Rising Event Index (0 - 65535)	5
Falling Event Index (0 - 65535)	0
Alarm Type	Rising
Sample Type	Delta
Polling Interval (1 - 65535)	60
Owner	Alarm_for_icmplnEchos

2. Click **Submit**.

3. Apply, verify, and save the configuration.



RMON group 9—events

The RMON Event group allows you to define events that are triggered by alarms. An event can be a log message, an SNMP trap message, or both.

When an alarm is generated, it triggers a corresponding event notification. Use the `/cfg/rmon/alarm x/revtidx` and `/fevtidx` commands to correlate an event index to an alarm.

RMON events use SNMP and syslogs to send notifications. Therefore, an SNMP trap host must be configured for trap event notification to work properly.

RMON uses a SYSLOG host to send syslog messages. Therefore, an existing SYSLOG host (`/cfg/sys/syslog`) must be configured for event log notification to work properly. Each log event generates a SYSLOG of type RMON that corresponds to the event.

Configuring RMON Events (CLI example)

1. Configure the RMON Event parameters.

```
>> /cfg/rmon/event 5 (Select RMON Event 5)
>> RMON Event 5# descrn "SYSLOG_generation_event"
>> RMON Event 5# type log
>> RMON Event 5# owner "Owner_event_5"
```

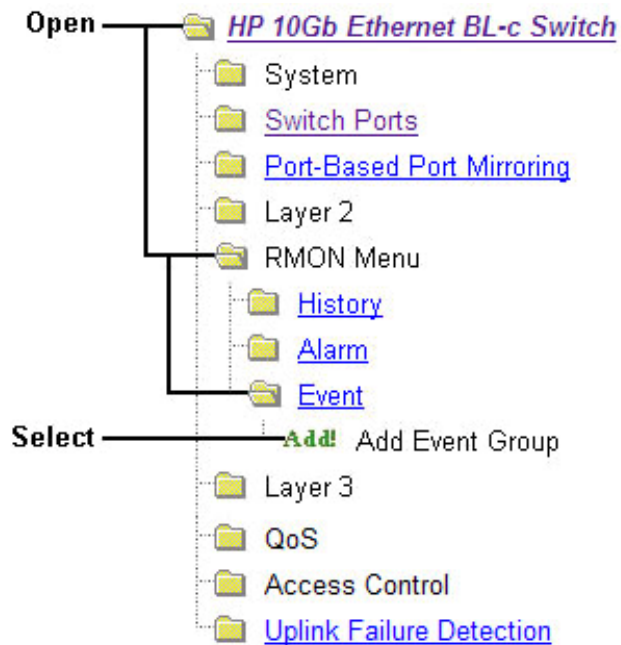
2. Apply and save the configuration.

```
>> RMON Alarm 5# apply (Make your changes active)
>> RMON Alarm 5# save (Save for restore after reboot)
```

This configuration creates an RMON event that sends a SYSLOG message each time it is triggered by an alarm.

Configuring RMON Events (BBI example)


1. Configure an RMON Event group.
 - a. Click the **Configure** context button.
 - b. Open the Switch folder, and select **RMON > Event > Add Event Group**.



- c. Configure RMON Event Group parameters. This configuration creates an RMON event that sends a SYSLOG message each time it is triggered by an alarm.

RMON Event Configuration	
Event Group ID (1 - 65535)	5
Event Type	Log
Description	SYSLOG_generation_event
Owner	Owner_event_5
<input type="button" value="Submit"/> <input type="button" value="Delete"/>	

2. Click **Submit**.
3. Apply, verify, and save the configuration.



High availability

Introduction

Switches support high availability network topologies. This release provides information about Uplink Failure Detection and Virtual Router Redundancy Protocol (VRRP).

Uplink Failure Detection

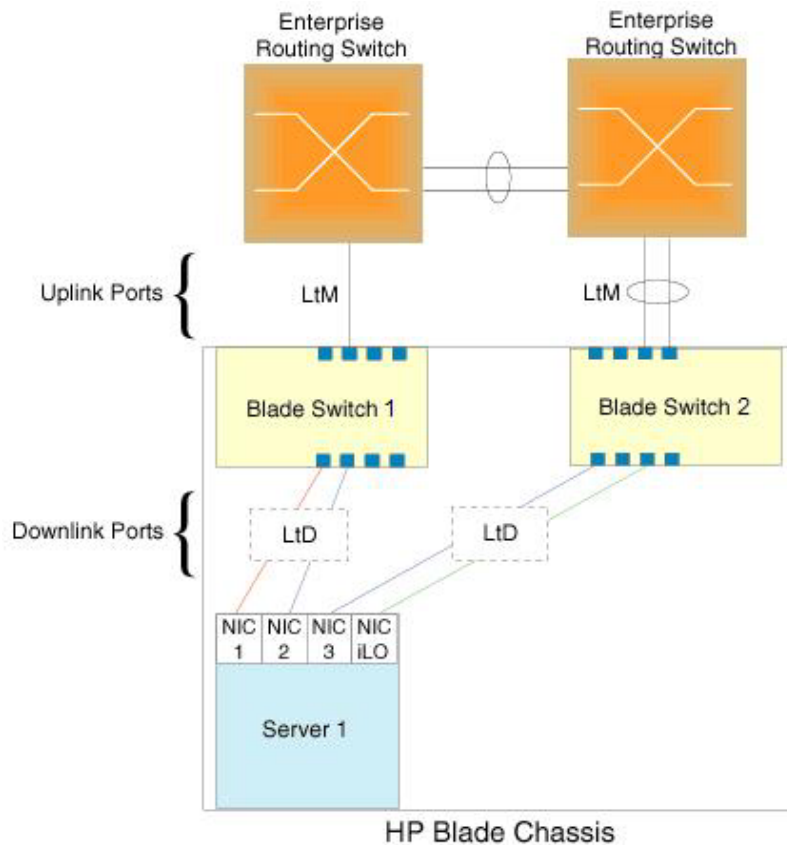
Uplink Failure Detection (UFD) is designed to support Network Adapter Teaming on HP server blades. For details about Network Adapter Teaming on HP ProLiant server blades, see the white paper at the following location:

<http://h18004.www1.hp.com/products/servers/networking/whitepapers.html>

UFD allows the switch to monitor specific uplink ports to detect link failures. When the switch detects a link failure, it automatically disables specific downlink ports. The corresponding server's network adapter can detect the disabled downlink, and trigger a network-adapter failover to another port on the switch, or another switch in the chassis.

The switch automatically enables the downlink ports when the uplink returns to service.

The following figure shows a basic UFD configuration, with a Failure Detection Pair (FDP) that consists of one LtM (Link to Monitor) and one LtD (Link to Disable). When the switch detects a link failure in the LtM, it disables the ports in the LtD. The server blade detects the disabled downlink port, which triggers a NIC failover.

Figure 24 Uplink Failure Detection for switches

Failure Detection Pair

To use UFD, you must configure a Failure Detection Pair and then turn UFD on. A Failure Detection Pair consists of the following groups of ports:

- Link to Monitor (LtM)

The Link to Monitor group consists of one uplink port (18-21), or one trunk group that contains only uplink ports. The trunk group can be a Link Aggregation Control Protocol trunk. The switch monitors the LtM for link failure.

- Link to Disable (LtD)

The Link to Disable group consists of one or more downlink ports (1-16) and trunk groups that contain only downlink ports. The trunk groups can be Link Aggregation Control Protocol trunks. When the switch detects a link failure on the LtM, it automatically disables all ports in the LtD.

When the LtM returns to service, the switch automatically enables all ports in the LtD.

Spanning Tree Protocol with UFD

If Spanning Tree Protocol (STP) is enabled on ports in the LtM, then the switch monitors the STP state and the link status on ports in the LtM. The switch automatically disables the ports in the LtD when it detects a link failure or STP Blocking state.

When the switch determines that ports in the LtM are in STP Forwarding State, then it automatically enables the ports in the LtD, to fall back to normal operation.

Configuration guidelines

This section provides important information about configuring UFD:

- UFD is required only when uplink-path redundancy is not available on the blade switches.
- Only one Failure Detection pair (one group of Links to Monitor and one group of Links to Disable) is supported on each switch (all VLANs and Spanning Tree Groups).
- An LtM can be either one uplink port or one Multi-Link trunk group of uplink ports.
- Ports that are already members of a trunk group are not allowed to be assigned to an LtM.
- A trunk group configured as an LtM can contain multiple uplink ports (18-21), but no downlink ports (1-16).
- An uplink port cannot be added to a trunk group if it already belongs to an LtM.
- An LtD can contain one or more ports, and/or one or more trunks
- A trunk group configured as an LtD can contain multiple downlink ports (1-16), but no uplink ports (18-21).

Monitoring Uplink Failure Detection

The UFD information menu displays the current status of the LtM and LtD, and their member ports or trunks. For example:

```
>> Information# ufd
Uplink Failure Detection: Enabled
LtM status: Down
Member      STG      STG State      Link Status
-----
port 19
           1      DISABLED
           10     DISABLED *
           15     DISABLED *
* = STP turned off for this port.

LtD status: Auto Disabled
Member      Link Status
-----
port 1      disabled
port 2      disabled
port 3      disabled
port 4      disabled
```

Use the `/stats/ufd` command to find out how many times link failure was detected on the LtM, how many times Spanning Tree blocking state was detected on the LtM, and how many times UFD disabled ports in the LtD.

Configuring Uplink Failure Detection

The preceding figure shows a basic UFD configuration. Assume that port 19 on Blade Switch 1 is connected to a Layer 2/3 routing switch outside of the chassis. Port 18 and port 19 on Blade Switch 2 form a trunk that is connected to a different Layer 2/3 routing switch.

In this example, NIC 1 is the primary network adapter; NIC 2, NIC 3, and NIC 4 are non-primary adapters. NIC 1 and NIC 2 are connected to port 1 and port 2 on Blade Switch 1. NIC 3 and NIC 4 are connected to port 1 and port 2 on Blade Switch 2.

Configuring UFD on Switch 1 (CLI example)

1. Assign uplink ports (18-21) to be monitored for communication failure.

```
>> Main# /cfg/ufd/fdp ena          (Enable Failure Detection Pair)
>> FDP# ltm                        (Select Link to Monitor menu)
>> Failure Link to Monitor# addport 19 (Monitor uplink port 19)
```

2. Assign downlink ports (1-16) to disable when an uplink failure occurs.

```
>> /cfg/ufd/fdp/ltd              (Select Link to Disable menu)
>> Failure Link to Disable# addport 1 (Add port 1 as a Link to Disable)
>> Failure Link to Disable# addport 2 (Add port 2 as a Link to Disable)
```

3. Turn UFD on.

```
>> /cfg/ufd/on                  (Turn Uplink Failure Detection on)
>> Uplink Failure Detection# apply (Make your changes active)
>> Uplink Failure Detection# save  (Save for restore after reboot)
```

When a link failure or Spanning Tree blocking occurs on port 19, Blade Switch 1 disables port 1 and port 2.

Configuring UFD on Switch 2 (CLI example)

1. Create a trunk group of uplink ports (18-21) to monitor.

```
>> Main# /cfg/trunk 2           (Create trunk group 2)
>> Trunk group 2# ena          (Enable trunk group 2)
>> Trunk group 2# add 18       (Add port 18 to trunk group 2)
>> Trunk group 2# add 19       (Add port 19 to trunk group 2)
```

2. Assign the trunk group to be monitored for communication failure.

```
>> Main# /cfg/ufd/fdp ena      (Enable Failover Pair)
>> FDP# ltm                    (Select Link to Monitor menu)
>> Failover Link to Monitor# addtrnk 2 (Monitor trunk group 2)
```

3. Assign downlink ports (1-16) to disable when an uplink failure occurs.

```
>> Main# /cfg/ufd/fdp/ltd      (Select Link to Disable menu)
>> Failover Link to Disable# addport 1 (Add port 1 as a Link to Disable)
>> Failover Link to Disable# addport 2 (Add port 2 as a Link to Disable)
```

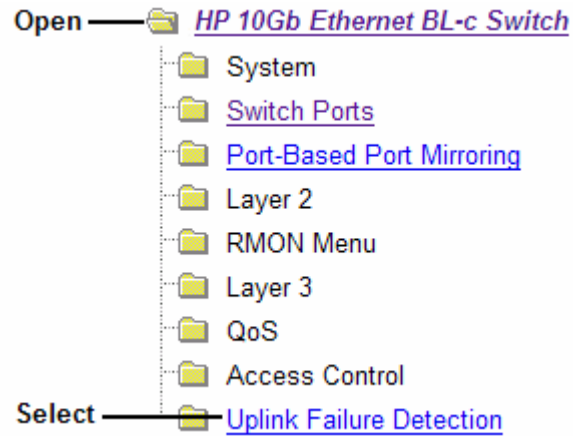
4. Turn UFD on.

```
>> Main# /cfg/ufd/on           (Turn Uplink Failure Detection on)
>> Uplink Failure Detection# apply (Make your changes active)
>> Uplink Failure Detection# save  (Save for restore after reboot)
```

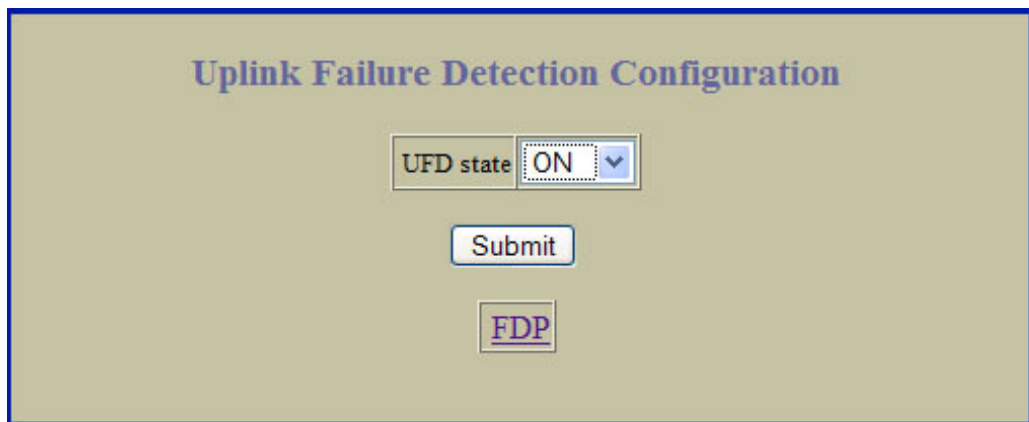
When a link failure or Spanning Tree blocking occurs on trunk group 2, Blade Switch 2 disables port 1 and port 2.

Configuring Uplink Failure Detection (BBI example)

1. Configure Uplink Failure Detection.
 - a. Click the **Configure** context button.
 - b. Open the Switch folder, and select **Uplink Failure Detection** (click the underlined text, not the folder).




- c. Turn Uplink Failure Detection on, and then select **FDP**.



- d. Enable the FDP. Select ports in the LtM Ports Available list, and click **Add** to place the ports into the Link to Monitor (LtM). Select ports in the LtD Ports Available list, and click **Add** to place the ports into the Link to Disable (LtD).

- e. Click **Submit**.
2. Apply, verify, and save the configuration.



VRRP overview

In a high-availability network topology, no device can create a single point-of-failure for the network or force a single point-of-failure to any other part of the network. This means that your network will remain in service despite the failure of any single device. To achieve this usually requires redundancy for all vital network components.

VRRP enables redundant router configurations within a LAN, providing alternate router paths for a host to eliminate single points-of-failure within a network. Each participating VRRP-capable routing device is configured with the same virtual router IP address and ID number. One of the virtual routers is elected as the master, based on a number of priority criteria, and assumes control of the shared virtual router IP address. If the master fails, one of the backup virtual routers will take control of the virtual router IP address and actively process traffic addressed to it.

With VRRP, Virtual Interface Routers (VIR) allows two VRRP routers to share an IP interface across the routers. VIRs provide a single Destination IP (DIP) for upstream routers to reach various servers, and provide a virtual default Gateway for the server blades.

VRRP components

Each physical router running VRRP is known as a VRRP router.

Virtual router

Two or more VRRP routers can be configured to form a virtual router (RFC 2338). Each VRRP router may participate in one or more virtual routers. Each virtual router consists of a user-configured virtual router identifier (VRID) and an IP address.

Virtual router MAC address

The VRID is used to build the virtual router MAC Address. The five highest-order octets of the virtual router MAC Address are the standard MAC prefix (00-00-5E-00-01) defined in RFC 2338. The VRID is used to form the lowest-order octet.

Owners and renters

Only one of the VRRP routers in a virtual router may be configured as the IP address owner. This router has the virtual router's IP address as its real interface address. This router responds to packets addressed to the virtual router's IP address for ICMP pings, TCP connections, and so on.

There is no requirement for any VRRP router to be the IP address owner. Most VRRP installations choose not to implement an IP address owner. For the purposes of this chapter, VRRP routers that are not the IP address owner are called renters.

Master and backup virtual router

Within each virtual router, one VRRP router is selected to be the virtual router master. See “Selecting the Master VRRP Router” for an explanation of the selection process.

NOTE: If the IP address owner is available, it will always become the virtual router master.

The virtual router master forwards packets sent to the virtual router. It also responds to Address Resolution Protocol (ARP) requests sent to the virtual router's IP address. Finally, the virtual router master sends out periodic advertisements to let other VRRP routers know it is alive and its priority.

Within a virtual router, the VRRP routers not selected to be the master are known as virtual router backups. Should the virtual router master fail, one of the virtual router backups becomes the master and assumes its responsibilities.

Virtual Interface Router

At Layer 3, a Virtual Interface Router (VIR) allows two VRRP routers to share an IP interface across the routers. VIRs provide a single Destination IP (DIP) for upstream routers to reach various destination networks, and provide a virtual default Gateway.

NOTE: Every VIR must be assigned to an IP interface, and every IP interface must be assigned to a VLAN. If no port in a VLAN has link up, the IP interface of that VLAN is down, and if the IP interface of a VIR is down, that VIR goes into INIT state.

VRRP operation

Only the virtual router master responds to ARP requests. Therefore, the upstream routers only forward packets destined to the master. The master also responds to ICMP ping requests. The backup does not forward any traffic, nor does it respond to ARP requests.

If the master is not available, the backup becomes the master and takes over responsibility for packet forwarding and responding to ARP requests.

Selecting the master VRRP router

Each VRRP router is configured with a priority between 1 and 254. A bidding process determines which VRRP router is or becomes the master—the VRRP router with the highest priority.

The master periodically sends advertisements to an IP multicast address. As long as the backups receive these advertisements, they remain in the backup state. If a backup does not receive an advertisement for three advertisement intervals, it initiates a bidding process to determine which VRRP router has the highest priority and takes over as master.

If, at any time, a backup determines that it has higher priority than the current master does, it can preempt the master and become the master itself, unless configured not to do so. In preemption, the backup assumes the role of master and begins to send its own advertisements. The current master sees that the backup has higher priority and will stop functioning as the master.

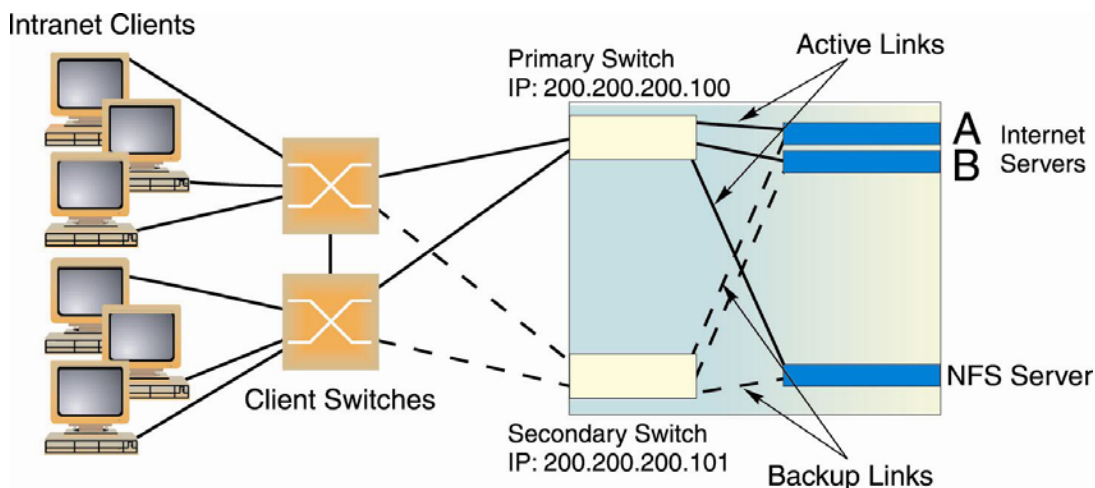
A backup router can stop receiving advertisements for one of two reasons—the master can be down, or all communications links between the master and the backup can be down. If the master has failed, it is clearly desirable for the backup (or one of the backups, if there is more than one) to become the master.

NOTE: If the master is healthy but communication between the master and the backup has failed, there will then be two masters within the virtual router. To prevent this from happening, configure redundant links to be used between the switches that form a virtual router.

Failover methods

With service availability becoming a major concern on the Internet, service providers are increasingly deploying Internet traffic control devices, such as application switches, in redundant configurations. Traditionally, these configurations have been hot-standby configurations, where one switch is active and the other is in a standby mode. A non-VRRP hot-standby configuration is shown in the figure below:

Figure 25 Non-VRRP hot-standby configuration



While hot-standby configurations increase site availability by removing single points-of-failure, service providers increasingly view them as an inefficient use of network resources because one functional application switch sits by idly until a failure calls it into action. Service providers now demand that vendor equipment support redundant configurations where all devices can process traffic when they are healthy, increasing site throughput and decreasing user response times when no device has failed.

The HP 10GbE switch high availability configurations are based on VRRP. The switch software implementation of VRRP includes proprietary extensions.


The switch software implementation of VRRP supports the Active-Active mode of high availability.

Active-Active redundancy

In an active-active configuration, shown in the following figure, two switches provide redundancy for each other, with both active at the same time. Each switch processes traffic on a different subnet. When a failure occurs, the remaining switch can process traffic on all subnets.

The following figure shows an Active-Active configuration example.

Figure 26 Active-Active redundancy



HP 10GbE switch extensions to VRRP

This section describes VRRP enhancements that are implemented in switch software:

Tracking VRRP router priority

The HP 10GbE switch software supports a tracking function that dynamically modifies the priority of a VRRP router, based on its current state. The objective of tracking is to have, whenever possible, the master bidding processes for various virtual routers in a LAN converge on the same switch. Tracking ensures that the selected switch is the one that offers optimal network performance. For tracking to have any effect on virtual router operation, preemption must be enabled.

Switch software can track the attributes listed in the following table:

Table 23 VRRP tracking parameters

Parameter	Description
Number of IP interfaces on the switch that are active (up) <code>/cfg/l3/vrrp/track/ifs</code>	Helps elect the virtual routers with the most available routes as the master. (An IP interface is considered active when there is at least one active port on the same VLAN.) This parameter influences the VRRP router's priority in virtual interface routers.
Number of active ports on the same VLAN <code>/cfg/l3/vrrp/track/ports</code>	Helps elect the virtual routers with the most available ports as the master. This parameter influences the VRRP router's priority in virtual interface routers.
Number of virtual routers in master mode on the switch <code>/cfg/l3/vrrp/track/vrs</code>	Useful for ensuring that traffic for any particular client/server pair is handled by the same switch, increasing routing efficiency. This parameter influences the VRRP router's priority in virtual interface routers.

Each tracked parameter has a user-configurable weight associated with it. As the count associated with each tracked item increases (or decreases), so does the VRRP router's priority, subject to the weighting associated with each tracked item. If the priority level of a standby is greater than that of the current master, then the standby can assume the role of the master.

See "Configuring the Switch for Tracking" for an example on how to configure the switch for tracking VRRP priority.

Virtual router deployment considerations

Review the following issues described in this section to prevent network problems when deploying virtual routers:

- Assigning VRRP Virtual Router ID
- Configuring the Switch for Tracking

Assigning VRRP virtual router ID

During the software upgrade process, VRRP virtual router IDs are assigned automatically if failover is enabled on the switch. When configuring virtual routers at any point after upgrade, virtual router ID numbers (`/cfg/13/vrrp/vr #/vrid`) must be assigned. The virtual router ID may be configured as any number between 1 and 250.

Configuring the switch for tracking

Tracking configuration largely depends on user preferences and network environment. Consider the configuration shown in the previous figure. Assume the following behavior on the network:

- Switch A is the master router upon initialization.
- If Switch A is the master and it has one fewer active servers than Switch B, then Switch A remains the master.
- This behavior is preferred because running one server down is less disruptive than bringing a new master online and severing all active connections in the process.
- If Switch A is the master and it has two or more active servers fewer than Switch B, then Switch B becomes the master.
- If Switch B is the master, it remains the master even if servers are restored on Switch A such that it has one fewer or an equal number of servers.
- If Switch B is the master and it has one active server fewer than Switch A, then Switch A becomes the master.

The user can implement this behavior by configuring the switch for tracking as follows:

1. Set the priority for Switch A to 101.
2. Leave the priority for Switch B at the default value of 100.
3. On both switches, enable tracking based on ports (`ports`), interfaces (`ifs`), or virtual routers (`vr`). You can choose any combination of tracking parameters, based on your network configuration.

NOTE: There is no shortcut to setting tracking parameters. The goals must first be set and the outcomes of various configurations and scenarios analyzed to find settings that meet the goals.


High availability configurations

The HP 10GbE switches offer flexibility in implementing redundant configurations. This section discusses the Active-Active configuration.

Active-Active configuration

The following figure shows an example configuration, where two switches are used as VRRP routers in an active-active configuration. In this configuration, both switches respond to packets.

Figure 27 Active-Active high availability configuration



Although this example shows only two switches, there is no limit on the number of switches used in a redundant configuration. It is possible to implement an active-active configuration across all the VRRP-capable switches in a LAN.

Each VRRP-capable switch in an active-active configuration is autonomous. Switches in a virtual router need not be identically configured.

In the scenario illustrated in the figure, traffic destined for IP address 10.0.1.1 is forwarded through the Layer 2 switch at the top of the drawing, and ingresses Switch A on port 20. Return traffic uses default gateway 1 (192.168.1.1). If the link between Switch A and the Layer 2 switch fails, Switch B becomes the Master because it has a higher priority. Traffic is forwarded to Switch B. Return traffic uses default gateway 2 (192.168.2.1), and is forwarded through the Layer 2 switch at the bottom of the drawing.

To implement the active-active example, perform the following switch configuration.

Task 1: Configure Switch A

1. Configure ports.

```

/cfg/l2/vlan 10           (Select VLAN 10)
>> VLAN 10# ena         (Enable VLAN 10)
>> VLAN 10# add 20      (Add port 20 to VLAN 10)
>> VLAN 10# ..
>> Layer 2# vlan 20     (Select VLAN 20)
>> VLAN 20# ena         (Enable VLAN 20)
>> VLAN 20# add 21      (Add port 21 to VLAN 20)

```

2. Configure client and server interfaces.

```

/cfg/l3/if 1 (Select interface 1)
>> IP Interface 1# addr 192.168.1.100 (Define IP address for interface 1)
>> IP Interface 1# vlan 10 (Assign VLAN 10 to interface 1)
>> IP Interface 1# ena (Enable interface 1)
>> IP Interface 1# ..
>> Layer 3# if 2 (Select interface 2)
>> IP Interface 2# addr 192.168.2.101 (Define IP address for interface 2)
>> IP Interface 1# vlan 20 (Assign VLAN 20 to interface 2)
>> IP Interface 2# ena (Enable interface 2)
>> IP Interface 2# ..
>> Layer 3# if 3 (Select interface 3)
>> IP Interface 3# addr 10.0.1.100 (Define IP address for interface 3)
>> IP Interface 3# mask 255.255.255.0 (Define subnet mask for interface 3)
>> IP Interface 3# ena (Enable interface 3)
>> IP Interface 2# ..
>> Layer 3# if 4 (Select interface 4)
>> IP Interface 4# addr 10.0.2.101 (Define IP address for interface 4)
>> IP Interface 4# mask 255.255.255.0 (Define subnet mask for interface 4)
>> IP Interface 4# ena (Enable interface 4)

```

3. Configure the default gateways. Each default gateway points to one of the Layer 2 routers.

```

/cfg/l3/gw 1 (Select default gateway 1)
>> Default gateway 1# addr 192.168.1.1 (Point gateway to the first L2 router)
>> Default gateway 1# ena (Enable the default gateway)
>> Default gateway 1# ..
>> Layer 3# gw 2 (Select default gateway 2)
>> Default gateway 1# addr 192.168.2.1 (Point gateway to the second router)
>> Default gateway 1# ena (Enable the default gateway)

```

4. Turn on VRRP and configure two Virtual Interface Routers.

```

/cfg/l3/vrrp/on (Turn VRRP on)
>> Virtual Router Redundancy Protocol# vr 1 (Select virtual router 1)
>> VRRP Virtual Router 1# vrid 1 (Set VRID to 1)
>> VRRP Virtual Router 1# if 1 (Set interface 1)
>> VRRP Virtual Router 1# addr 192.168.1.200 (Define IP address)
>> VRRP Virtual Router 1# ena (Enable virtual router 1)
>> VRRP Virtual Router 1# .. (Enable virtual router 1)
>> Virtual Router Redundancy Protocol# vr 2 (Select virtual router 2)
>> VRRP Virtual Router 2# vrid 2 (Set VRID to 2)
>> VRRP Virtual Router 2# if 2 (Set interface 2)
>> VRRP Virtual Router 2# addr 192.168.2.200 (Define IP address)
>> VRRP Virtual Router 2# ena (Enable virtual router 2)

```

5. Enable tracking on ports. Set the priority of Virtual Router 1 to 101, so that it becomes the Master.

```

/cfg/l3/vrrp/vr 1 (Select VRRP virtual router 1)
>> VRRP Virtual Router 1# track/ports/ena (Set tracking on ports)
>> VRRP Virtual Router 1 Priority Tracking# ..
>> VRRP Virtual Router 1# prio 101 (Set the VRRP priority)
>> VRRP Virtual Router 1# ..
>> Virtual Router Redundancy Protocol# vr 2 (Select VRRP virtual router 2)
>> VRRP Virtual Router 1# track/ports/ena (Set tracking on ports)

```

6. Turn off Spanning Tree Protocol globally.

```

/cfg/l2/stg 1/off (Turn off STG)
>> Spanning Tree Group 1# apply
>> Spanning Tree Group 1# save

```

Task 2: Configure Switch B

1. Configure ports.

```

/cfg/l2/vlan 10                (Select VLAN 10)
>> VLAN 10# ena                (Enable VLAN 10)
>> VLAN 10# add 20             (Add port 20 to VLAN 10)
>> VLAN 10# ..
>> Layer 2# vlan 20           (Select VLAN 20)
>> VLAN 20# ena                (Enable VLAN 20)
>> VLAN 20# add 21            (Add port 21 to VLAN 20)

```

2. Configure client and server interfaces.

```

/cfg/l3/if 1                    (Select interface 1)
>> IP Interface 1# addr 192.168.1.101 (Define IP address for interface 1)
>> IP Interface 1# vlan 10          (Assign VLAN 10 to interface 1)
>> IP Interface 1# ena              (Enable interface 1)
>> IP Interface 1# ..
>> Layer 3# if 2                  (Select interface 2)
>> IP Interface 2# addr 192.168.2.100 (Define IP address for interface 2)
>> IP Interface 1# vlan 20          (Assign VLAN 20 to interface 2)
>> IP Interface 2# ena              (Enable interface 2)
>> IP Interface 2# ..
>> Layer 3# if 3                  (Select interface 3)
>> IP Interface 3# addr 10.0.1.101   (Define IP address for interface 3)
>> IP Interface 3# mask 255.255.255.0 (Define subnet mask for interface 3)
>> IP Interface 3# ena              (Enable interface 3)
>> IP Interface 2# ..
>> Layer 3# if 4                  (Select interface 4)
>> IP Interface 4# addr 10.0.2.100   (Define IP address for interface 4)
>> IP Interface 4# mask 255.255.255.0 (Define subnet mask for interface 4)
>> IP Interface 4# ena              (Enable interface 4)

```

3. Configure the default gateways. Each default gateway points to one of the Layer 2 routers.

```

/cfg/l3/gw 1                    (Select default gateway 1)
>> Default gateway 1# addr 192.168.2.1 (Point gateway to the first L2 router)
>> Default gateway 1# ena              (Enable the default gateway)
>> Default gateway 1# ..
>> Layer 3# gw 2                  (Select default gateway 2)
>> Default gateway 1# addr 192.168.1.1 (Point gateway to the second router)
>> Default gateway 1# ena              (Enable the default gateway)

```

4. Turn on VRRP and configure two Virtual Interface Routers.

```

/cfg/l3/vrrp/on                 (Turn VRRP on)
>> Virtual Router Redundancy Protocol# vr 1 (Select virtual router 1)
>> VRRP Virtual Router 1# vrid 1          (Set VRID to 1)
>> VRRP Virtual Router 1# if 1           (Set interface 1)
>> VRRP Virtual Router 1# addr 192.168.1.200 (Define IP address)
>> VRRP Virtual Router 1# ena              (Enable virtual router 1)
>> VRRP Virtual Router 1# ..              (Enable virtual router 1)
>> Virtual Router Redundancy Protocol# vr 2 (Select virtual router 2)
>> VRRP Virtual Router 2# vrid 2          (Set VRID to 2)
>> VRRP Virtual Router 2# if 2           (Set interface 2)
>> VRRP Virtual Router 2# addr 192.168.2.200 (Define IP address)
>> VRRP Virtual Router 2# ena              (Enable virtual router 2)

```

5. Enable tracking on ports. Set the priority of Virtual Router 2 to 101, so that it becomes the Master.

```

/cfg/l3/vrrp/vr 1 (Select VRRP virtual router 1)
>> VRRP Virtual Router 1# track/ports/ena (Set tracking on ports)
>> VRRP Virtual Router 1 Priority Tracking# ..
>> VRRP Virtual Router 1# ..
>> Virtual Router Redundancy Protocol# vr 2 (Select VRRP virtual router 2)
>> VRRP Virtual Router 2# track/ports/ena (Set tracking on ports)
>> VRRP Virtual Router 2 Priority Tracking# ..
>> VRRP Virtual Router 2# prio 101 (Set the VRRP priority)

```

6. Turn off Spanning Tree Protocol globally. Apply and save changes.

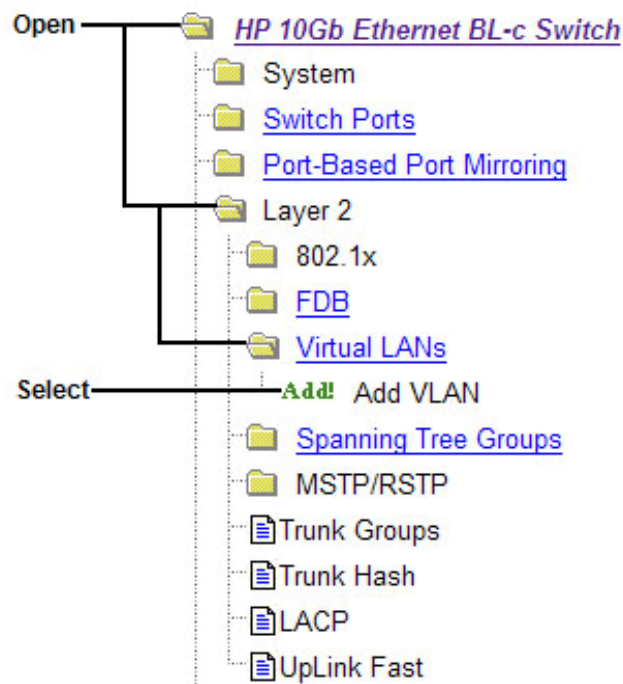
```

/cfg/l2/stg 1/off (Turn off STG)
>> Spanning Tree Group 1# apply
>> Spanning Tree Group 1# save

```

Task 1: Configure Switch A (BBI example)

1. Configure ports and VLANs.
 - a. Click the **Configure** context button.
 - b. Open the Virtual LANs folder, and select **Add VLAN**.



- c. Configure port 20 as a member of VLAN 10 and port 21 as a member of VLAN 20. Enable each VLAN.

VLAN "New" Configuration

VLAN Name	VLAN Ten
VLAN ID (1 - 4095) From	10
VLAN State	enabled ▼
Spanning Tree Group	1

Ports Available

PortID	▲
Port:1	
Port:2	
Port:3	
Port:4	
Port:5	
Port:6	
Port:7	
Port:8	
Port:9	▼

Ports in Vlan

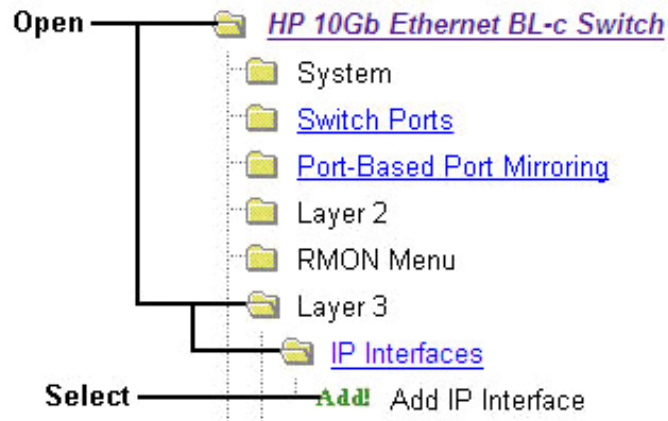
PortID	▲
Port:20	
	▼

- d. Click **Submit**.

2. Configure the following client and server interfaces:

- IF 1
IP address = 192.168.1.100
Subnet mask = 255.255.255.0
VLAN 10
- IF 2
IP address = 10.10.12.1
Subnet mask = 255.255.255.0
VLAN 20
- IF 3
IP address = 10.10.12.1
Subnet mask = 255.255.255.0
- IF 4
IP address = 10.10.12.1
Subnet mask = 255.255.255.0

- a. Open the IP Interfaces folder, and select **Add IP Interface**.




- b. Configure an IP interface. Enter the IP address, subnet mask, and VLAN membership. Enable the interface.

IP Interface Configuration

IP Interface Identifier (1-250)	<input type="text" value="1"/>
IP Address	<input type="text" value="192.169.1.100"/>
Enabled?	<input type="text" value="Enabled"/> ▾
Subnet Mask	<input type="text" value="255.255.255.0"/>
VLAN Membership ID (1 - 4095)	<input type="text" value="10"/>
Enable Disable BOOTP Relay	<input type="text" value="Enabled"/> ▾

- c. Click **Submit**.

3. Configure the default gateways. Each default gateway points to one of the Layer 2 routers.
 - a. Open the Default Gateways folder, and select **Add Default Gateway**.



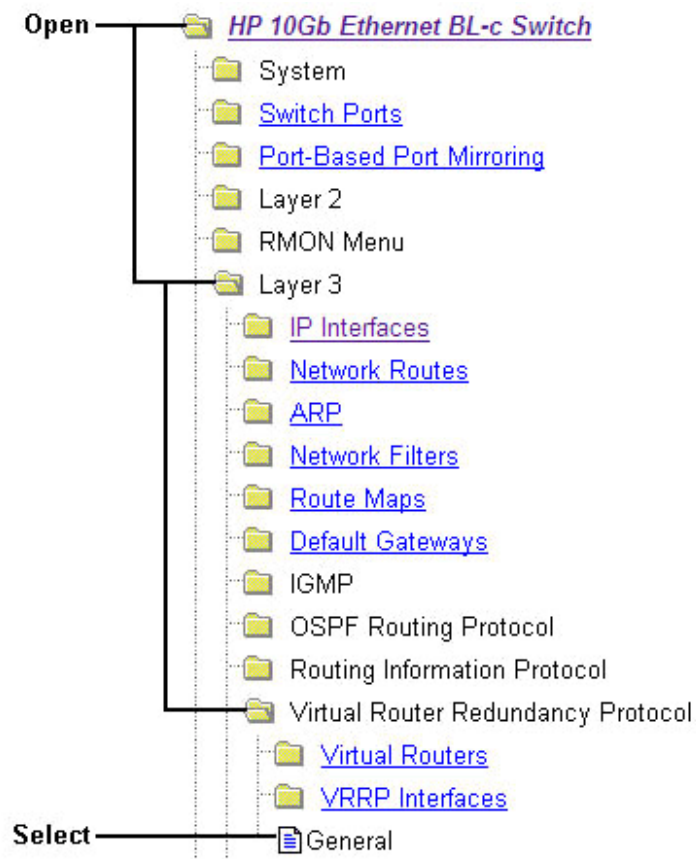
- b. Configure the IP address for each default gateway. Enable the default gateways.

Default Gateway Configuration

Default Gateway Identifier 1-4, 254	<input type="text" value="1"/>
Default Gateway IP Address	<input type="text" value="192.168.1.1"/>
Enable/Disable Default Gateway	<input type="text" value="Enabled"/> ▼
Enable/Disable ARP only health checks	<input type="text" value="Disabled"/> ▼
Health Check Interval (0-60 sec)	<input type="text" value="2"/>
Retries before Out of Service (1-120)	<input type="text" value="8"/>

- c. Click **Submit**.

4. Turn on VRRP and configure two Virtual Interface routers.
 - a. Open the Virtual Router Redundancy Protocol folder, and select **General**.



- b. Enable VRRP processing.

VRRP General Configuration


VRRP Processing Enabled?

VRRP virtual router tracking increment (0-254)	<input type="text" value="2"/>
VRRP IP interface tracking increment (0-254)	<input type="text" value="2"/>
VRRP VLAN switch port tracking increment (0-254)	<input type="text" value="2"/>

VRRP Virtual Router Group Configuration

Virtual Router Identifier (1- 250)	<input type="text" value="1"/>	IP interface (1- 249)	<input type="text" value="1"/>
Enabled?	<input type="text" value="Disabled"/>	Priority (1- 254)	<input type="text" value="100"/>
Advertisement Interval (1- 255)	<input type="text" value="1"/>	Owner Preemption?	<input type="text" value="Enabled"/>
Track other IP interfaces?	<input type="text" value="Disabled"/>	Track VLAN switch ports?	<input type="text" value="Disabled"/>

- c. Click **Submit**.
- d. Open the Virtual Routers folder, and select **Add Virtual Router**.




- e. Configure the IP address for Virtual Router 1 (VR1). Enable tracking on ports, and set the priority to 101. Enable The Virtual Router.

Virtual Router Configuration

Virtual Router Number (1- 250)	<input type="text" value="1"/>
Virtual Router Identifier (1- 250)	<input type="text" value="1"/>
IP Address	<input type="text" value="192.168.1.200"/>
IP interface (1-249)	<input type="text" value="1"/>
Enabled?	Enabled <input type="button" value="v"/>
Priority (1- 254)	<input type="text" value="101"/>
Advertisement Interval (1- 255)	<input type="text" value="1"/>
Owner Preemption?	Enabled <input type="button" value="v"/>
Track master virtual routers?	Disabled <input type="button" value="v"/>
Track other IP interfaces?	Disabled <input type="button" value="v"/>
Track VLAN switch ports?	Enabled <input type="button" value="v"/>

- f. Click **Submit**.
- g. Select **Add Virtual Router**.




- h. Configure the IP address for Virtual Router 2 (VR2). Enable tracking on ports, but set the priority to 100 (default value). Enable The Virtual Router.

Virtual Router Configuration

Virtual Router Number (1- 250)	<input type="text" value="2"/>
Virtual Router Identifier (1- 250)	<input type="text" value="2"/>
IP Address	<input type="text" value="192.168.2.200"/>
IP interface (1-249)	<input type="text" value="2"/>
Enabled?	Enabled <input type="button" value="v"/>
Priority (1- 254)	<input type="text" value="100"/>
Advertisement Interval (1- 255)	<input type="text" value="1"/>
Owner Preemption?	Enabled <input type="button" value="v"/>
Track master virtual routers?	Disabled <input type="button" value="v"/>
Track other IP interfaces?	Disabled <input type="button" value="v"/>
Track VLAN switch ports?	Enabled <input type="button" value="v"/>

- i. Click **Submit**.
5. Turn off Spanning Tree globally.
- a. Open the Spanning Tree Groups folder, and select **Add Spanning Tree Group**.



- b. Select a Spanning Tree Group.

Switch Spanning Tree Groups Configuration

1. Search Range
Spanning Trees(1-128) From To

2. Search Options
Bridge Priority(0 = any) State
Search Operation

Spanning Tree Group	State	Bridge Priority	Bridge Hello Time	Bridge Max Age	Bridge Forward Delay	Bridge Aging Time
1	on	32768	2	20	15	300
2	on	32768	2	20	15	300
3	on	32768	2	20	15	300
4	on	32768	2	20	15	300
5	on	32768	2	20	15	300
6	on	32768	2	20	15	300
7	on	32768	2	20	15	300
8	on	32768	2	20	15	300

Select →

- c. Enter Spanning Tree Group ID 1 and set the Switch Spanning Tree State to **off**.

Switch Spanning Tree Group Configuration

Spanning Tree Group ID (1-128)	<input type="text" value="1"/>
Switch Spanning Tree State	<input type="text" value="off"/> ▼
Bridge Priority (0-65535)	<input type="text" value="32768"/>
Bridge Hello Time (1-10secs)	<input type="text" value="2"/>
Bridge Max Age (6-40secs)	<input type="text" value="20"/>
Bridge Forward Delay (4-30secs)	<input type="text" value="15"/>

VLANs Available

Vlan ID:Name
 2:VLAN Name
 4095:Mgmt VLAN

VLANs in STG

Vlan ID:Name
 1:Default VLAN

Switch Spanning Tree Port Configuration

Switch Port	Port Priority	Port Path Cost	Port Spanning Tree State
<u>1</u>	128	2	off
<u>2</u>	128	2	off

- d. Click **Submit**.

6. Apply, verify, and save the configuration.



Troubleshooting tools

Introduction

This appendix discusses some tools to help you use the Port Mirroring feature to troubleshoot common network problems on the switch.


Port Mirroring

The Port Mirroring feature on the switch is very useful for troubleshooting any connection-oriented problem. Any traffic in or out of one or more ports can be mirrored to a single monitoring port to which a network monitor can be attached.

Port Mirroring can be used as a troubleshooting tool or to enhance the security of your network. For example, an Intrusion Detection Service (IDS) server can be connected to the monitor port to detect intruders attacking the network.

As shown in the following figure, port 18 is monitoring ingress traffic (traffic entering the switch) on port 21 and egress traffic (traffic leaving the switch) on port 1. You can attach a device to port 18 to monitor the traffic on ports 21 and 1.

Figure 28 Port Mirroring



This figure shows two mirrored ports monitored by a single port. Similarly, you can have one mirrored port to one monitored port, or many mirrored ports to one monitored port. The switch does not support a single port being monitored by multiple ports because it supports only one monitored port configured at a time.

Ingress traffic is duplicated and sent to the mirrored port before processing, and egress traffic is duplicated and sent to the mirrored port after processing.

Configuring Port Mirroring (CLI example)

To configure Port Mirroring for the example shown in the preceding figure:

1. Specify the monitoring port.

```
>> # /cfg/pmirr/monport 18           (Select port 18 for monitoring)
```

2. Select the ports that you want to mirror.

```
>> Port 18 # add 21                   (Select port 21 to mirror)
>> Enter port mirror direction [in, out, or both]: in (Monitor ingress traffic
                                                    on port 21)
>> Port 18 # add 1                   (Select port 1 to mirror)
>> Enter port mirror direction [in, out, or both]: out (Monitor egress traffic
                                                    on port 1)
```

3. Enable Port Mirroring.

```
>> # /cfg/pmirr/mirr ena             (Enable port mirroring)
```

4. Apply and save the configuration.

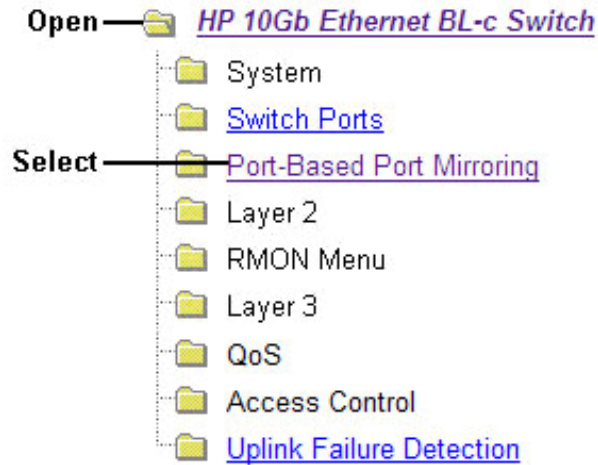
```
>> PortMirroring# apply              (Apply the configuration)
>> PortMirroring# save               (Save the configuration)
```

5. View the current configuration.

```
>> PortMirroring# cur                (Display the current settings)
Port mirroring is enabled
Monitoring Ports Mirrored Ports
1      none
2      none
3      none
4      none
5      none
      :
      :
16     none
18 (21, in) (1, out)
20     none
21     none
```


Configuring Port Mirroring (BBI example)

1. Configure Port Mirroring.
 - a. Click the **Configure** context button.
 - b. Open the Switch folder, and select **Port-Based Port Mirroring** (click the underlined text, not the folder).



- c. Click a port number to select a monitoring port.

Port-Based Port Mirroring Configuration

Enable Port-Based Port Mirroring? ▼

Port Mirroring Table


Monitoring Port	Mirrored Ports
<u>1</u>	none
<u>2</u>	none
<u>3</u>	none
<u>4</u>	none
<u>5</u>	none
<u>6</u>	none
<u>7</u>	none
<u>8</u>	none
<u>9</u>	none
<u>10</u>	none
<u>11</u>	none
<u>12</u>	none
<u>13</u>	none
<u>14</u>	none
<u>15</u>	none
<u>16</u>	none
<u>18</u>	none
<u>19</u>	none
<u>20</u>	none
<u>21</u>	none

Select —

- d. Click **Add Mirrored Port**.

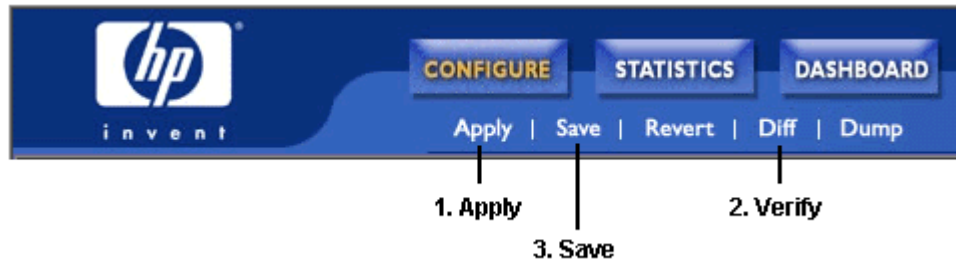


- e. Enter a port number for the mirrored port, and select the Port Mirror Direction.




- f. Click **Submit**.

2. Apply, verify, and save the configuration.



3. Verify the Port Mirroring configuration on the switch.



Other network troubleshooting techniques

Other network troubleshooting techniques include the following.

Console and Syslog messages

When a switch experiences a problem, review the console and Syslog messages. The switch displays these informative messages when state changes and system problems occur. Syslog messages can be viewed by using the `/info/sys/log` command. For more information on interpreting syslog messages, see the *HP 10Gb Ethernet BL-c Switch Command Reference Guide*.

Ping

To verify station-to-station connectivity across the network, execute the following command:

```
ping <host name> | <IP address> [ (attempts, 1-32) [ msec delay ]] [-m|-mgt|-d|-data]
```

The IP address is the hostname or IP address of the device. The number of attempts is optional.

Msec delay (optional) is the number of milliseconds between attempts. By default, the `-m` or `-mgt` option for the management port is used. To use data ports, specify the `-d` or `-data` option.

Trace route

To identify the route used for station-to-station connectivity across the network, execute the following command:

```
tracert <host name> | <IP address> [<max-hops> [ msec delay ]]
```

The IP address is the hostname or IP address of the target station. Max-hops (optional) is the maximum distance to trace (1-32 devices). Msec delay (optional) is the number of milliseconds to wait for the response.

Statistics and state information

The switch keeps track of a large number of statistics and many of these are error condition counters. The statistics and state information can be very useful when troubleshooting a LAN or Real Server problem.

For more information about available statistics, see one of the following:

- "Viewing statistics" chapter of the *HP 10Gb Ethernet BL-c Switch Browser-based Interface Reference Guide*, or
- "Statistics Menu" chapter of the *HP 10Gb Ethernet BL-c Switch Command Reference*

Customer support tools

The following diagnostics tools are not user-configurable and should be performed through HP technical support.

- **Offline Diagnostics**—This tool is used for troubleshooting suspected switch hardware issues. These tests verify that the selected hardware is performing within expected engineering specifications.
- **Software Panics**—If a fatal software condition is found during runtime, the switch will capture the current hardware and software state information into a panic dump. This dump file can be analyzed post-mortem to determine the cause of the problem.

- Stack Trace—If a fatal software condition occurs, the switch dumps stack trace data to the console. If you have a console attached to the switch, capture the console dump, and forward it to HP technical support.

Index

- 8**
 - 802.1x port states, 49
- A**
 - accessing the switch: defining
 - source IP addresses, 20;
 - RADIUS authentication, 21;
 - security, 20; using the command line interface (CLI), 12
 - ACL Blocks and Groups, 90
 - ACL configuration examples, 92
 - ACL filters, 87
 - active-active redundancy, 175
 - allowable source IP addresses, 20
- B**
 - BBI: See Browser-Based Interface, 134
 - Bridge Protocol Data Unit (BPDU), 68
 - broadcast domains, 53, 58
 - Browser-Based Interface, 134
- C**
 - Cisco EtherChannel, 38
 - Command Line Interface, 134
 - command line interface (CLI), 12
 - configuration, 38
 - configuration rules: port mirroring, 38; trunking, 38; VLANs, 38
 - configuring: OSPF, 141
 - configuring SNMP trap hosts, 17
 - console messages, 195
 - customer support tools, 195
- D**
 - default gateway: configuration example, 13, 109
 - Differentiated Services Code Point (DSCP)DSCP, 97
 - Dynamic Host Configuration Protocol, 112
- E**
 - EAPoL configuration guidelines, 51
 - EtherChannel, as used with port trunking, 38
 - Extensible Authentication Protocol over LAN (EAPoL)EAPoL, 47
- F**
 - switch failover: overview, 175
 - Fast Uplink Convergence, 76
- fault tolerance, port trunking, 37
- FDB static entries, 66
- frame tagging, 55
- G**
 - default gateway, 108
- H**
 - high-availability, 9, 167
- I**
 - IBM DirectorSNMP: IBM Director, 14
 - IEEE 802.1Q VLAN tagging, 55, 58
 - IEEE standards: 802.1x, 47
 - IGMP Filtering, configuring (BBI example), 125
 - IGMP Filtering, configuring (CLI example), 122
 - IGMP Snooping, 118
 - IGMP Snooping, configuring (BBI example), 123
 - IGMP Snooping, configuring (CLI example), 121
 - IGMPv3: snooping, 119
 - IP address: routing example, 109
 - IP interfaces: configuring, 12; example configuration, IP interface configuration, 111; VLAN #1 (default), 58; VLANs, 58
 - IP routing, 106; default gateway configuration, 13, 109; IP interface configuration, 109, 111; IP interfaces example configuration, 109
 - IP subnets: routing; IP subnets, 107; routing, 108
 - IP subnets, VLANs, 53
- J**
 - fragmenting jumbo frames: fragmenting to normal size; routing, 106
- L**
 - LACP, 44
 - Link Aggregation Control Protocol, 44
 - logical segment, 53
 - LSAs, 133
- M**
 - Main Menu, command line interface (CLI), 12
 - media access control (MAC) address, 12
 - meter, 91
 - mirroring ports, 191
 - monitoring ports, 191
 - MSTP, 77
 - multi-links between switches, using port trunking, 36
 - multiple spanning tree groups, 71
- N**
 - network management, 14
- O**
 - Open Shortest Path First (OSPF), 131
 - ICMP;IGMP;TCP;UDP: filtering criteria, 87
 - OSPF: overview, 131
 - OSPF: area types, 131
 - OSPF: neighbors, 133
 - OSPF: link state database, 133
 - OSPF: route summarization, 137
 - OSPF: default route; OSPF, 137
 - OSPF: virtual link, 138
 - OSPF: authentication, 139
 - OSPF: host routes; OSPF, 140
 - OSPF: external routes, 141
 - OSPF configuration examples, 141, 142, 150, 151, 152
 - OSPF configuration, verifying, 154
- P**
 - ping, 195
 - Port Fast Forwarding, 76
 - port mirroring: configuration rules, 38; troubleshooting, 191
 - port trunking, 37
 - port trunking, fault tolerance, 37
 - Port-based Network Access control, 47
 - Port-based traffic control, 51
 - ports: monitoring, 191; physical, 54
 - priority value (802.1p), 98
 - PVID (port VLAN ID), 54

Q

Quality of Service, 86
 queuing and scheduling, 105

R

RADIUS: port 1812 and 1645, 88;
 port 1813, 88
 redundancy: active-active, 175;
 VRRP (Virtual Router
 Redundancy Protocol), 175
 re-mark, 91
 Remote Authentication Dial-in User
 Service (RADIUS):
 authentication, 21; SSH/SCP, 34
 Remote monitoring (RMON), 155
 RIP (Routing Information Protocol):
 advertisements, 114; distance
 vector protocol, 114; hop count,
 114; metric, 114
 RIP configuration, example, 117
 RIP features, 115
 RMON (remote monitoring), 155
 RMON groups: alarms, 161; events,
 165; examples of, 156, 159, 160,
 161, 162, 164, 165, 166, 170;
 history, 158; statistics, 155
 Router ID: OSPF; router ID, 138
 routers, 107, 109; border;
 peer;autonomous systems (AS),
 134; switch-based routing
 topology, 108
 routing: internal and
 external;internal routing, 134

RSA keys, 33

RSTP, 77

S

security, 20; RADIUS
 authentication, 21; switch
 management, 20; VLANs, 53
 segmentation, 53
 HP-OpenView: HP-OpenView, 14
 SNMP, 14
 SNMP, 134
 SNMP v1.0, 14
 SNMP v3.0, 14
 Spanning Tree Protocol (STP),
 multiple instances, 72
 spanning tree, configuration rules,
 38
 SSH: RSA host and server keys,
 33; supported clients, 31
 SSH/SCP configuring, 31
 Static MRouter, configuring (CLI
 example), 122
 Static Multicast Router, configuring
 (BBI example), 129
 statistics, 195
 switch management: security, 20;
 via IP interface, 58
 switch ports VLANs membership,
 54
 syslog messages, 195

T

tagging, 55
 technical terms: port VLAN identifier
 (PVID), 55; tagged frame, 55;
 tagged member, 55; untagged
 frame, 55; untagged member,
 55; VLAN identifier (VID), 55
 Telnet, 11
 trace route, 195
 Trunk Hash algorithm, 44
 trunking: configuration rules, 38
 typographical conventions, 10

V

VLANs: broadcast domains, 53, 58,
 110; configuration rules, 38;
 default PVID, 54; example
 showing multiple VLANs, 60; ID
 numbers, 53; IP interface
 configuration, 111; IP interfaces,
 58; multiple spanning trees, 68;
 multiple VLANs, 55; port
 members, 54; PVID, 54; security,
 53; Spanning-Tree Protocol, 68;
 tagging, 54; topologies, 58;
 VLAN #1 (default), 58
 VRRP (Virtual Router Redundancy
 Protocol): active-active
 redundancy, 175; overview, 173;
 virtual router ID numbering; ID
 numbering, 177; overview;
 tracking, 176; vrid, 173